

Approximate Bayesian Computation for Astronomical Model Analysis: A Case Study in Galaxy Demographics and Morphological Transformation at High Redshift

E. Cameron^{1*} and A. N. Pettitt¹

¹*School of Mathematical Sciences (Statistical Science), Queensland University of Technology (QUT), GPO Box 2434, Brisbane 4001, QLD, Australia*

Submitted to MNRAS: 7 Feb 2012.

ABSTRACT

“Approximate Bayesian Computation” (ABC) represents a powerful methodology for the analysis of complex stochastic systems for which the likelihood of the observed data under an arbitrary set of input parameters may be entirely *intractable*—the latter condition rendering useless the standard machinery of tractable likelihood-based, Bayesian statistical inference (e.g. conventional Markov Chain Monte Carlo simulation; MCMC). In this article we demonstrate the potential of ABC for astronomical model analysis by application to a case study in the morphological transformation of high redshift galaxies. To this end we develop, first, a stochastic model for the competing processes of merging and secular evolution in the early Universe; and second, through an ABC-based comparison against the observed demographics of the first generation of massive ($M_{\text{gal}} > 10^{11} M_{\odot}$) galaxies (at $1.5 < z < 3$) in the CANDELS/EGS dataset we derive posterior probability densities for the key parameters of this model. The “Sequential Monte Carlo” (SMC) implementation of ABC exhibited herein, featuring both a self-generating target sequence and self-refining MCMC kernel, is amongst the most efficient of contemporary approaches to this important statistical algorithm. We highlight as well through our chosen case study the value of careful summary statistic selection, and demonstrate two modern strategies for assessment and optimisation in this regard. Ultimately, our ABC analysis of the high redshift morphological mix returns tight constraints on the evolving merger rate in the early Universe and reveals merging, rather than secular evolution, as the most important mechanism for building up the first generation of bulges in early-type disks.

Key words:

Galaxies: evolution – galaxies: formation – methods: statistical.

1 INTRODUCTION

With origins in population genetics and evolutionary biology (e.g. Tavaré et al. 1997; Pritchard et al. 1999; Beaumont et al. 2002; and see Csilléry et al. 2010 for a recent review) Approximate Bayesian Computation (ABC) offers a powerful technique for recovering posterior probability densities from complex stochastic models for which the likelihood of the observed data under a given set of input parameters may be entirely intractable. Examples include the estimation of time to the most recent common ancestor under the coalescent model with recombination given a full suite of modern DNA sequencing (Marjoram & Tavaré 2006), or the derivation of transition probabilities in contin-

uous time Markov models of macroparasite population evolution from simple demographics (Drovandi & Pettitt 2010). However, although there exist a variety of important *astro-physical* models with inherently intractable likelihoods, to-date applications of ABC in this field remain surprisingly rare.¹ The only indispensable ingredients required for ABC

¹ Indeed the authors can find no astronomical reference to either the terms “approximate Bayesian computation” or “likelihood-free” (inference) on the NASA ADS database; and Google Scholar indicates no astronomical citations yet to any of the biological/mathematical ABC literature mentioned herein. The power-point slides from a talk by Chad Schafer at the Statistical Challenges in Modern Astronomy V conference in 2011 available at URL(<http://astrostatistics.psu.edu/su11scma5/scma5.html>) detailing two potential uses for ABC in extra-galactic data analysis represent to our knowledge the only prior application in this field.

* E-mail: dr.ewan.cameron@gmail.com

are: (i) a *stochastic model* for the observed data, replicating the behaviour of all random processes driving the system at hand as well as any relevant observational errors; and (ii) a *discrepancy measure*, based typically on a set of low order summary statistics, to quantitatively gauge similarity between output from this model and the empirical benchmark.

One potentially valuable role for ABC in an astronomical context *may* thus be in the constraint of semi-analytic models (SAMs) of galaxy formation (cf. Cole et al. 2000; Benson et al. 2003; Baugh 2006; Bower et al. 2006; De Lucia et al. 2010; Neistein & Weinmann 2010)—in which the output at run-time necessarily exhibits complex stochasticity owing to the effects of cosmic variance (induced computationally via sampling from within large-scale dark matter simulations, Springel et al. 2001; Knebe et al. 2011; or via Monte Carlo construction of halo merger trees, Lacey & Cole 1993; Parkinson et al. 2008). For ABC analysis of such codes an appropriate discrepancy parameter might then be the metric distance between simulated and observed luminosity functions under a sensible binning scheme.² In the (perhaps-not-so) disparate fields of Active Galactic Nucleus (AGN) power spectrum analysis (Uttley et al. 2002; Ingram & Done 2011; Kelly et al. 2011), asteroseismology (Kjeldsen & Bedding 2004; Hekker et al. 2011), and initial mass function profiling (Fumagalli et al. 2011; Da Silva et al. 2011) one also routinely encounters models from which random simulations can readily be made, but for which the corresponding likelihood functions (featuring correlated noise or population clustering) are inherently intractable. Where conditions (i) and (ii) above are thus satisfied ABC offers an easily-implemented, theoretically well-established (Nunes & Balding 2010; Marin et al. 2011; Fearnhead & Prangle 2012) alternative to the computationally intensive “approximate likelihood” approach (requiring very large scale simulation/re-simulation, e.g. Wood 2010; Henriques et al. 2009; Lu et al. 2011; and note Benson et al. 2011 regarding the required diversity of merger trees sampled for genuine convergence of SAMs), or the user-intensive application of model emulators (requiring a non-trivial degree of run-time supervision and operator expertise, cf. Bower et al. 2010 and references therein).³

² For readers familiar with the work of Bower et al. (2010) we note that the “discrepancy parameter” introduced for their emulation of the GALFORM SAM could not be employed as such in ABC as it is not in fact a gauge of model–data similarity; indeed it serves an entirely different purpose in their analysis, acting as an error term for cosmic variance and structural uncertainty in their code.

³ As a caveat to the above referencing we note that: (i) though the analyses of Henriques et al. (2009) and Lu et al. (2011) are both conducted broadly in the style of the “approximate likelihood” approach formalised by Wood (2010) their specific implementations contain certain ad hoc features of dubious statistical merit (e.g., the neglect of cosmic variance uncertainty and correlation between bins when computing the luminosity function likelihood in the former, and in the latter the arbitrary three-fold inflation of the observational error bars to account for uncertainty in these uncertainties); and (ii) though the work of Kampakoglou et al. (2008) has in previous papers been cited as an example of MCMC-based SAM constraint, in fact, their study concerns a purely analytic model for which there exists no in-

In this paper we illustrate heuristically the power of ABC for astronomical model analysis through application to a case study in the morphological transformation of massive galaxies at high redshift. In particular, we demonstrate a contemporary Sequential Monte Carlo (SMC) formulation of the ABC algorithm (cf. Del Moral et al. 2006; Sisson et al. 2007; Drovandi & Pettitt 2010), as well as a regression-based procedure for constructing an optimal summary statistic–discrepancy measure pairing for the purpose of parameter estimation (Fearnhead & Prangle 2012). Importantly, the stochastic model we explore herein features both an “independent evolution” case for which the likelihood is in fact tractable and a “co-evolution” case for which it is not—the former allowing the reliability of our ABC-based solution to be established against conventional Markov Chain Monte Carlo (MCMC) simulation and the latter a demonstration of the unique possibilities of ABC analysis.

Installation of the new Wide-Field Camera 3 (WFC3) on the Hubble Space Telescope (HST) in 2009—and the subsequent allocation of vast amounts of observing time to deep, near-infrared (NIR) surveys with this instrument, including the Early Release Science program (ERS; Windhorst et al. 2011) and the Cosmic Assembly Near-IR Deep Extragalactic Legacy Survey (CANDELS; Grogin et al. 2011; Koekemoer et al. 2011)—has at last made accessible (at high resolution) the rest-frame optical morphologies of distant galaxies at the epoch of peak cosmic star formation and AGN activity ($z \sim 2$; Lilly et al. 1996; Madau et al. 1996; Oesch et al. 2012; Warren et al. 1994). Early studies exploiting these new datasets have documented the emergence of the first Hubble sequence analogues (Cameron et al. 2011b; Conselice et al. 2011a; Szomoru et al. 2011a), demonstrated the compactness of the first massive spheroids (Szomoru et al. 2010, 2011b; Newman et al. 2011), explored the unique characteristics of galaxies ultraluminous at infrared (Kartaltepe et al. 2011) and X-ray wavelengths (Kocevski et al. 2012; Rosario et al. 2011; Schawinski et al. 2011), and probed structural transformation in extreme cluster environments (Lotz et al. 2011; Papovich et al. 2011). Thus far, however, there have been remarkably few studies to exploit the full potential of *demographic analysis* for constraining pathways of galaxy evolution—one early exemplar being Bell et al.’s (2011) search for correlations between the global observables of key galaxy sub-populations in the CANDELS dataset divided coarsely by rest-frame optical morphological type (via the usual proxy of global Sérsic index; cf. Driver et al. 2006; Cameron & Driver 2009; Kelvin et al. 2011). Hence, we have chosen here specifically for our exposition of the ABC technique a case study in the demographic analysis of WFC3 data in the hope of motivating further research in this direction.

intrinsic stochasticity (thus, only approximate observational errors enter their likelihood computation). Finally, we refer the interested reader to Hartig et al. (2011) for a concise overview of the similarities and differences between the “approximate likelihood” and ABC approaches to inference from statistical simulation, and to Nott et al. (2011) for an advanced treatment of the link between a particular version of ABC and the Bayes Linear technique (cf. Goldstein & Wooff 2007) underlying the model emulator approach.

The structure of this paper is as follows. In Section 2 we review the publicly-available source catalogues and images comprising our high redshift demographic benchmark, then in Section 3 we present the core of our case study in ABC for astronomical model analysis. First, we describe our model for galaxy evolution and our procedure for stochastic simulation from this model (Section 3.1). Second, we explain the ABC algorithm and the SMC approach to its implementation (Section 3.2). Third, we examine in depth the important process of constructing an optimal summary statistic–discrepancy parameter pairing (Section 3.3). And fourth, we confirm the similarity between our ABC and MCMC posteriors in the tractable “independent evolution” case, and present our final ABC-only posteriors for the more realistic, but likelihood intractable, “co-evolution” case (Section 3.4). In Section 4 we conclude this paper with a discussion of the implications of the model constraints so derived for astrophysical theories of morphological transformation in the early Universe.

We have thus attempted to organise our exposition of ABC in such a manner as to allow astronomers interested in this important statistical algorithm but not working directly in the area of galaxy evolution to optionally skip over the technical details and justification of our model (Section 3.1) without disadvantage (instead reading only Sections 3.2, 3.3, and 3.4 in depth). All magnitudes are quoted in the AB system and a standard $\{\Omega_M = 0.3, \Omega_\Lambda = 0.7, h = 0.7\}$ cosmological model is adopted throughout.

2 DATA

Featuring a vast ensemble of multiwavelength imaging compiled from both ground-based and space-based observatories the Extended Groth Strip (EGS) region of the Northern Sky (centred on RA: $14^h 17^m$, Dec: $+52^\circ 30'$) numbers amongst the premier legacy survey fields of the modern era. The All-wavelength Extended Groth strip International Survey team (AEGIS; Davis et al. 2007) has been responsible for the bulk of this data collection through extensive observational campaigns with HST and Spitzer. Such a comprehensive set of photometric measurements greatly facilitates the estimation of redshifts and stellar masses via Spectral Energy Distribution (SED) template fitting, and there exist a number of published studies characterising the high redshift galaxy population in the EGS to this effect.

In this study we employ the publicly-available and up-to-date, ultra-violet-(UV)-to-far-infrared-(FIR)-based catalogue of Barro et al. (2011a,b) [B11 hereafter] to identify a complete sample of high mass ($M_{\text{gal}} > 10^{11} M_\odot$), early Universe ($1.5 < z < 3$) systems. The B11 photometric redshifts, based on up to 19 band flux measurements in the survey core, feature an overall accuracy (measured against a spectroscopic subsample from AGEIS with median $z \sim 1.3$) of $\Delta z/(1+z) = 0.034$ at a sub-2% catastrophic failure rate. At the highest redshifts ($z > 2.5$) comparison against a spectroscopic sample of 91 Lyman-Break Galaxies (LBGs) confirms only a slight degradation to $\Delta z/(1+z) = 0.069$. The corresponding B11 stellar masses for these systems were derived

using the PEGASE⁴ SED library (with Salpeter initial mass function and Calzetti extinction)—the choice of which (from amongst the wide range of alternative SED libraries) represents the dominant source of systematic uncertainty here (of order 0.1-0.3 dex; Barro et al. 2011b). For the purposes of this paper, in which our principle aim is to demonstrate as straightforwardly as possible the technicalities of the ABC approach, we hereafter neglect further quantitative consideration of these uncertainties (except when required for fitting the build up in number density over cosmic time, which contributes two “nuisance parameters” to our model, in Section 3.1).

The CANDELS team (Grogin et al. 2011; Koekemoer et al. 2011) is currently engaged in the acquisition of high-resolution, near-infrared (and UV) imaging targeting distant galaxies in selected sub-regions of five key legacy fields (GOODS-N, GOODS-S⁵, the EGS, COSMOS⁶, and the UDS⁷) totaling $\sim 800 \text{ arcmin}^2$ under an allocation of 902 orbits of HST/WFC3 exposure time. Drizzled to a pixel scale of 0.06 arcsec, the presently-available epoch (egs01) of imaging within (an $\sim 90 \text{ arcmin}^2$ sub-region of) the EGS field features a point spread function (PSF) full width half maximum (FWHM) of ~ 0.18 arcsec and a 5σ detection limit of 26.8 mag in the F160W (*H*-band) filter (with comparable coverage in the F125W filter). As such CANDELS already represents the highest quality dataset published to-date for the study of rest-frame optical morphologies at $z \sim 2-3$ in the EGS. Accordingly for the present analysis we derive our high redshift demographic benchmark from visual classification of all 126 members of the B11 catalogue at $M_{\text{gal}} > 10^{11} M_\odot$ and $1.5 < z < 3$ imaged thus far.

To this end one of us (EC) inspected each source carefully in the CANDELS (HST WFC3/IR) *H*-band mosaic with ds9 and assigned it one of the following four types: (i) *spheroid* (compact elliptical; cf. Szomoru et al. 2011b), (ii) *spheroid-plus-disk* (early-type disk with a prominent central bulge; cf. Cameron et al. 2011b), (iii) *pure disk* (late-type, bulgeless [perhaps clumpy] disk; cf. Elmegreen et al. 2007a), or (iv) *ongoing merger* (evident violent relaxation event in progress, as revealed by the presence of distinctive tidal features and/or multiple massive nuclei; cf. Elmegreen et al. 2007b). In total we count 8 Type I spheroids, 9 Type II spheroid-plus-disks, 90 Type III pure disks, and 19 Type IV ongoing mergers in our sample. Example *H*-band postage stamp images characterising these archetypal high redshift morphologies are presented in the lefthand panel of Figure 1, as is an illustration of the demographic evolution across our sample in the righthand panel. Once again in accordance with the expository aims of this paper regarding ABC we do not explore the (complex) possible impacts of classification subjectivity on our results—although we note that both ABC and the Bayesian framework in general offer a powerful statistical basis for marginalising over such sources of uncer-

⁴ PEGASE: Projet d’Etude des GALaxies par Synthese Evolutive (Fioc & Rocca-Volmerange 1997).

⁵ GOODS-N/-S: Great Observatories Origins Deep -North/-South.

⁶ COSMOS: COSMic evOLution Survey.

⁷ UDS: UKIDSS (UKIRT [United Kingdom InfraRed Telescope] Infrared Deep Sky Survey) ultra-Deep Survey.

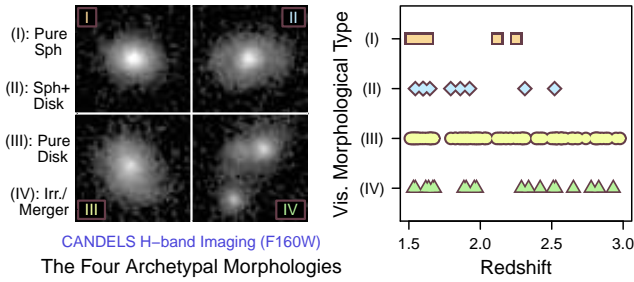


Figure 1. (Left:) CANDELS (HST WFC3/IR) H -band postage stamp images characterising the four archetypal morphologies present amongst massive ($M_{\text{gal}} > 10^{11} M_{\odot}$), high redshift ($1.5 < z < 3$) galaxies in the B11 dataset. (Right:) An illustration of the evolving morphological mix amongst our B11 (CANDELS/EGS) sample.

tainty (cf. Gelman et al. 2003; E.N. Taylor, in prep.; and see our treatment of various “nuisance parameters” in Sections 3.1 and 3.2), particularly where the experimental evaluation of the classification system has been appropriately designed and implemented (Hand 1997). Reassuringly though, the relative proportions of early and late type systems recovered from our classification process are at least broadly consistent with those reported by Buitrago et al. (2011) in their analysis of the (lower resolution) GOODS NICMOS Survey (GNS, Conselice et al. 2011b; also Mortlock et al. 2011).

3 STATISTICAL METHODOLOGY & RESULTS

Here we begin by introducing our stochastic model for the morphological transformation of high redshift galaxies, detailing both the tractable “independent evolution” case and the intractable “co-evolution” case, in Section 3.1. We then proceed to outline the SMC approach to ABC in Section 3.2, and to demonstrate linear regression-based construction of an optimal summary statistic–discrepancy parameter pairing for our model in Section 3.3. Finally, in Section 3.4 we compare the performance of SMC ABC against “tractable likelihood”-based MCMC in the “independent evolution” case and present our ABC-only solution for the more realistic “co-evolution” case.

3.1 Morphological Transformation as a Stochastic Process

With the current generation of SAMs yet to offer detailed or reliable predictions for the morphologies of simulated galaxies (Almeida et al. 2007; González et al. 2009) we develop here instead a basic stochastic model for describing the competing processes of morphological transformation in the early Universe. In this endeavour we are motivated both by contemporary observational results and hydrodynamical simulations. The purpose of simulation in this study is thus *not* to work forwards through parameterised approximations for the physical laws of halo accretion, gas cooling, and star formation (amongst others) in order to constrain their “fundamental” scaling coefficients (as in SAMs), but rather to explore in a rigorous statistical sense the extent to which the rates of incidence of the key events thought to shape

morphological evolution are jointly constrained by the observed demographics. Nevertheless, working backwards from these constraints (our posterior probability densities) one may hopefully achieve insight into the underlying physical mechanisms, as we discuss in Section 4.

As a starting point for our model we suppose that the arrival of galaxies at the top end of the high redshift stellar mass function may be faithfully represented as a non-homogeneous Poisson birth process with an underlying rate, $\lambda_b(t)$, increasing as $10^K t^\gamma$ from a zero baseline at $z = 6$. Thus, on average (i.e., over an infinite volume), $\Phi_{>10^{11} M_{\odot}}(z) = \int_0^z \lambda_b(t) dt = 10^K \frac{t^{\gamma+1}}{\gamma+1}$ (modulo the impact of merging amongst $M_{\text{gal}} > 10^{11} M_{\odot}$ systems, which is intrinsically rare at these redshifts, i.e., negligible in this context; cf. Man et al. 2012 and our discussion in Section 4.1). In the lefthand panel of Figure 2 we illustrate the build up in number density at $M_{\text{gal}} > 10^{11} M_{\odot}$ over the interval $z \sim 1.5$ to 6 synthesised from observations in the GNS (Mortlock et al. 2011, M11), the MOIRCS⁸ Deep Survey (Kajisawa et al. 2009, K09), the NEWFIRM⁹ Medium-Band Survey (Marchesini et al. 2010, M10; Brammer et al. 2011, BR11), the UDS (Caputi et al. 2011, C11), and the EGS (the present study, C12). Interestingly, where their redshift baselines overlap a number of these rival $\Phi_{>10^{11} M_{\odot}}(z)$ determinations exhibit surprisingly large discrepancies, even with regard to their respective cosmic variance uncertainties (marked as 1σ error bars in Figure 2 following the recipe of Moster et al. 2011). As highlighted by Brammer et al. (2011) such discrepancies may well arise from the systematic errors inherent in SED-based stellar mass computation (owing to degeneracies between the various template libraries), and we suspect this to be the case here.

To estimate our birth rate parameters, K and γ , we thus perform standard MCMC exploration of the relevant posterior probability density space under a likelihood model in which the datapoints from each of the above-listed studies are assumed subject to a common systematic bias in addition to cosmic variance. The prior magnitude of this systematic bias component (in dex) is treated as normally-distributed with $\mathcal{N}(0, 0.1^2)$ for each survey. Our respective priors on K and γ are both Uniform, with the former uninformative and the latter standard (i.e., bound between zero and one). The joint posterior density for $\{K, \gamma\}$ thus recovered is roughly bivariate Normal with $f_{K, \gamma} \sim \mathcal{N}_{\text{Trunc.}}([-4.1, 0.65]', [0.06^2, 0.1^2; \rho = 0.05]; 0 < \gamma < 1)$. For reference we plot the corresponding (pointwise) median, 1σ , and 3σ credible intervals for $\Phi_{>10^{11} M_{\odot}}(z)$ against the various empirical determinations shown in Figure 2. Due to the relatively small cosmic volume probed by the CANDELS/EGS dataset we do not attempt to further constrain K and γ during our ABC analysis; instead we treat these two variables as “nuisance parameters” of our stochastic model and integrate them out at run-time (see Section 3.2).

It is important to note at this point that the marginal posterior density on the systematic bias in our EGS datapoints favours a (median) of $+0.11$ dex, suggesting that the B11 stellar masses are systematically *over*-estimated by a

⁸ MOIRCS: Multi-Object InfraRed Camera and Spectrograph.

⁹ NEWFIRM: NOAO (National Optical Astronomical Observatory) Extremely Wide-Field InfraRed iMager.

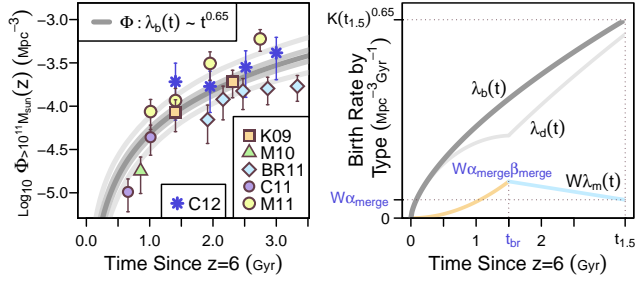


Figure 2. (Left:) The build up in galaxy number density at stellar masses $M_{\text{gal}} > 10^{11} M_{\odot}$ from $z = 6$ to $z = 1.5$ (in terms of cosmological time since $z = 6$) synthesised from recent observational determinations of $\Phi_{>10^{11} M_{\odot}}(z)$ from the literature. The K09, M10, C11, and M11 datapoints shown here are derived from integration of their published Schechter mass function fits, while the BR11 datapoints are sourced directly from that paper. The error bar on each indicates the 1σ contribution of cosmic variance for the respective survey and bin width (following the recipe of Moster et al. 2011). The dark, medium, and light grey curves plotted underneath represent our (pointwise) median, 1σ , and 3σ credible intervals, respectively, on the evolving mean number density given the posteriors of the fitted K and γ for our non-homogeneous Poisson birth process model, $\lambda_b(t) = 10^K t^\gamma$. The median curve corresponds roughly to $\gamma \sim 0.65$. (Right:) Illustration of the birth rate by type in our model (cf. Section 3.1). The rate at which sub- $10^{11} M_{\odot}$ galaxies are promoted above this mass threshold by merging is taken as W times our $M_{\text{gal}} > 10^{11} M_{\odot}$ merger rate, $W\lambda_m(t)$; leaving the rate of promotion (by star formation) of Type III disks, $\lambda_d(t)$, as the remainder with respect to the total birth rate, $\lambda_b(t)$.

corresponding ~ 0.10 dex (adopting the $z \sim 1.5$ mass function slope of Mortlock et al. 2011). Hence, it is perhaps more appropriate to describe our B11 CANDELS/EGS dataset as an $M_{\text{gal}} \gtrsim 10^{11} M_{\odot}$ sample, stressing the inherent uncertainty in SED-based stellar mass selection.

We next suppose that each galaxy arrives at the top end of the high redshift stellar mass function as either a (star-forming) late-type disk (Type III) or an ongoing major merger (Type IV)—a simplifying assumption which serves to reduce markedly the required dimensionality of our model, yet which is also consistent with the present state of knowledge on this topic. In a recent empirical census of rest-frame optical morphology amongst the sub- $10^{11} M_{\odot}$ population at $1.5 < z < 3.5$ Cameron et al. (2011b) were unable to identify a single unambiguous spheroid beyond $z \approx 2.2$ in their sample from the ERS (and see also Conselice et al. 2011a for a similar result). Amongst the small fraction ($\sim 20\%$) of sub- $10^{11} M_{\odot}$ spheroids discovered in their sample at later epochs only one was found to be actively star-forming—leaving dry merging as perhaps the only feasible (but also unlikely, cf. Lin et al. 2010; Chou et al. 2011) mechanism for sub- $10^{11} M_{\odot}$ spheroids to thus move above this threshold mass without transitioning through a standard Type IV phase. Meanwhile, contemporary hydrodynamical simulations have demonstrated the theoretical potential for high redshift disks at $10^{10.5} - 10^{11} M_{\odot}$ to sustain immense rates of star formation fueled by cold flow gas accretion (Dekel et al. 2009; Brooks et al. 2009; Genel et al. 2010) while avoiding secular bulge assembly through the wind-driven disruption of clump instabilities (Genel et al. 2010; Hopkins et al.

2011)—ensuring their rapid transition to the high mass regime intact as Type III systems.

The probability of birth as a Type IV ongoing merger is estimated in our model as W times the ratio of the instantaneous merger rate amongst our $M_{\text{gal}} > 10^{11} M_{\odot}$ population, $\lambda_m(t_{\text{birth}})$, to the corresponding instantaneous birth rate, $\lambda_b(t_{\text{birth}})$, as shown in the righthand panel of Figure 2. This factor, W , represents another nuisance parameter of our model corresponding to the ratio by number density of galaxies in such a mass range that a single major merger could promote them above $10^{11} M_{\odot}$ to those already beyond this threshold. (There is an implicit assumption here that the merger rate does not evolve significantly with mass over this small baseline.) According to the shape of the stellar mass function at these redshifts (e.g. Brammer et al. 2011; Mortlock et al. 2011) we estimate $W \approx 0.5 \pm 0.2$. Important to note is the fact that the merger rate so defined, $\lambda_m(t)$, is strictly that of galaxies already at the top end of the stellar mass function—i.e., we are effectively eliminating the contribution to the observed merger fraction from galaxies that were sub- $10^{11} M_{\odot}$ prior to their most recent encounter. With $\lambda_b(t)$ the total birth rate and $W\lambda_m(t)$ the birth rate of Type IV mergers the corresponding birth rate of Type III disks, $\lambda_d(t)$, is simply the arithmetic difference, $\lambda_b(t) - W\lambda_m(t)$ (as indicated in the righthand panel of Figure 2).

As for the total birth rate described earlier, $\lambda_b(t)$, merging in our model is characterised as a non-homogeneous Poisson process, with a unique rate by volume of

$$\lambda_m(t) = \begin{cases} \frac{\alpha_{\text{merge}} \beta_{\text{merge}}}{t_{\text{br}}^2} t^2 & \text{for } 0 \leq t \leq t_{\text{br}}, \\ \frac{\alpha_{\text{merge}} \beta_{\text{merge}} - (t - t_{\text{br}}) \alpha_{\text{merge}} (\beta_{\text{merge}} - 1)}{t_{1.5} - t_{\text{br}}} & \text{for } t_{\text{br}} < t \leq t_{1.5}. \end{cases}$$

Here α_{merge} represents the baseline merger rate (in units of $\text{Mpc}^{-3} \text{Gyr}^{-1}$) at our lowest redshift, $z = 1.5$, with $\alpha_{\text{merge}} \beta_{\text{merge}}$ the peak cosmological merger rate for massive galaxies at t_{br} . This latter model parameter, t_{br} , thus dictates a point of phase transition (or “break”) beyond which the merger rate by volume must ultimately decrease (with *increasing redshift*) back to zero at $z = 6$ (the time origin of our model) at least as fast as the total number density of galaxies itself, lest the specific merger rate (per galaxy), $\frac{\lambda_m(t)}{\lambda_b(t)}$, become asymptotically infinite. Here we have chosen for simplicity a fixed, marginally sufficient decay rate for $\lambda_m(t)$ above this transition redshift of $t^{2 > \gamma(\approx 0.65)+1}$. One possible extension of our model, which may well be worthwhile in the future if/when a larger demographic dataset for $z \sim 1.5$ CANDELS sources becomes available, would be to treat this pre- t_{br} decay rate as a free parameter of the fit.

Previous empirical studies have argued alternately that the massive galaxy merger rate is either very near constant (de Ravel et al. 2011; Williams et al. 2011; Lotz et al. 2011; Man et al. 2012) or markedly increasing with redshift (Conselice et al. 2003, 2009a; López-Sanjuan et al. 2009; de Ravel et al. 2009; Bluck et al. 2009, 2011) from the local volume out to $z \sim 1.5$; and even less consensus exists regarding its behaviour to $z \sim 3$ and beyond (Bluck et al. 2011; Man et al. 2012; Law et al. 2012; Williams et al. 2011). The intrinsic clumpiness of galaxy-scale star-formation at the observed optical (i.e., rest-frame UV) wavelengths of many of these studies (most non-WFC3) has proved a persistent source of uncertainty, introducing substantial ambiguity into

the interpretation of those morphological signatures otherwise indicative of recent merging locally (cf. Conselice et al. 2003; López-Sanjuan et al. 2009). Uncertainties concerning the fraction of apparent close pairs to ultimately merge (Lotz et al. 2008) and the visibility timescales of the resulting post-merger tidal features (Lotz et al. 2010) have only compounded these difficulties. (Important to note is that the full demographic analysis performed in the present study permits a simultaneous, non-degenerate constraint of the latter unknown, which is a significant advantage of this particular mode of analysis.)

Our only inflexible constraints on the tuneable parameters of $\lambda_m(t)$ here are thus that α_{merge} is, of course, strictly positive, β_{merge} is greater than or equal to one, and $t_6(=0) \leq t_{\text{br}} \leq t_{1.5}$. Hence we adopt only weak priors specified as: (i) a T distribution in \log_{10} -space for α_{merge} [with $\mu = -4$, $\Sigma = 0.5$, and 10 degrees of freedom, truncated to a key region of interest at a lower bound of -5.5 and an upper bound of -2.5]; (ii) a Beta distribution in \log_{10} -space for $2\beta_{\text{merge}}$ [with shape coefficients, 1 and 4, favouring a smaller peak-to-baseline ratio over a higher one]; and (iii) a Beta distribution for t_{br} (as a fraction of $t_{1.5}$) [with shape coefficients, 2 and 1, favouring a break closer to $z \sim 1.5$ than $z \sim 6$]. The grey-shaded tiles and histograms in Figures 4, 9, and 10 offer graphical representations of these prior densities.

Since, as mentioned earlier, dry-merging appears to be remarkably uncommon in the early Universe—cf. the rapidly declining fraction of red-red pairs with increasing redshift (Lin et al. 2010; Chou et al. 2011; Kampczyk et al. 2011) and the overall paucity of passive galaxies, in general, above $z \sim 2$ (Brammer et al. 2011; Whitaker et al. 2011; Wuyts et al. 2011)—we assume that *all* mergers to occur under our model are gas-rich and therefore generate a distinctive post-merger morphology with irregular tidal features (Elmegreen et al. 2007b; Lotz et al. 2010). We model the (observed *H*-band) visibility timescale of these features according to a Gamma distribution with scale, $1 + 100\tau_{\text{irr morph}}$, and shape coefficient, 100, for $\tau_{\text{irr morph}}$ in Gyr; thereby allowing an ~ 0.1 Gyr interquartile spread to account for some intrinsic variation in the cold gas fraction (and thus the merger-to-stable-remnant transition time; Lotz et al. 2010) across the galaxy population sampled. Inspired by contemporary hydrodynamical simulations of gas-rich mergers (Lotz et al. 2010) we choose our prior density on $\tau_{\text{irr morph}}$ to favour timescales on the order of 0.2 to 0.7 Gyr (though permitting, at much lower prior density, the possibility of even >1 Gyr timescales; Conselice 2009b) by adopting a Beta distribution with shape coefficients, 3 and 5, on this parameter divided by 1.5. Upon fading of these post-merger tidal features we suppose the final remnant may assume either a Type I (pure spheroid) or Type II (spheroid-plus-disk) morphology with the probability of the latter outcome a tuneable parameter, $P_{\text{Sph+D remnant}}$. Given the ongoing debate within the hydrodynamical modelling community regarding the relative frequency of mergers conducive to disk reformation at these epochs (e.g. Robertson et al. 2006; Bournaud et al. 2011) we adopt a Beta distribution prior on $P_{\text{Sph+D remnant}}$ with shape coefficients, 1 and 3, favouring Type II production in less than one in every four mergers.

The second pathway of morphological transformation permitted under our model is that of secular evolution

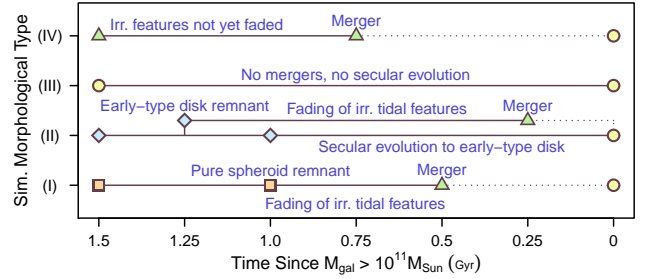


Figure 3. Schematic illustration of the five characteristic pathways of high redshift morphological transformation permitted under the stochastic model described in this paper (shown over an arbitrary timeline of 1.5 Gyr from “birth” to “observation” with an assumed Type III birth class). Solid lines are used to mark the fundamental pathways necessary to reach a given evolutionary state, whereas dashed lines allow for a range of possible degenerate evolutionary histories prior to the most recent merger (the details of which are inconsequential to the final state achieved). The secular evolution pathway to Type II status and the null evolution pathway to Type III status are the only branches for which one could not substitute Type IV as the birth class here.

of Type III disks to Type II, the theoretical mechanism proposed to drive this process being the inwards migration of massive star-forming clumps as encountered in certain hydrodynamical simulations (Bournaud et al. 2007; Elmegreen et al. 2008; Dekel et al. 2009). We again model this process stochastically via a Gamma distribution with scale, $1 + 50\tau_{\text{sec ev}}$, and shape coefficient, 50, for $\tau_{\text{sec ev}}$ in Gyr; thereby inducing an intrinsic spread in this variable at run-time intended to mimic the impact of natural diversity in the structure and kinematics of high redshift disks. Though inwards migration has been well publicised as the favoured hypothesis of the SINS¹⁰ team to explain the characteristic morphologies and SEDs of the clump population hosted amongst members of their pioneering $z \sim 2$ survey (Förster-Schreiber et al. 2011; Genzel et al. 2011), as noted earlier the most recent hydrodynamical simulations incorporating the effects of wind-driven mass loss, at least in the $10^{10.5}$ – $10^{11} M_{\odot}$ regime, indicate that a large fraction of these clumps may be too short-lived to migrate successfully into a central bulge (Genel et al. 2010; Hopkins et al. 2011). We therefore adopt such a prior on the timescale for secular bulge formation as to allow a full range of scenarios from rapid growth on sub-Gyr timescales (implying that many of our Type III disks should transition to Type II before $z \sim 1.5$) to incredibly slow growth on up to 10 Gyr timescales (implying none should). Mathematically we represent our prior density on this parameter via a Uniform distribution in \log_{10} -space bounded between -1 and 1 .

Figure 3 illustrates schematically the five distinct pathways of morphological transformation permitted under the above-specified model. We note for reference that Figures 4, 9, and 10 offer graphical representations of the prior densities on all our model parameters.

¹⁰ SINS: Spectroscopic Imaging survey in the Near-infrared with SINFONI [Spectrograph for INtegral Field Observations in the Near Infrared].

3.1.1 Simulation from our Stochastic Model

Having outlined above the principle details of our stochastic model for high redshift morphological transformation we now describe our corresponding procedure for simulating from this model under two distinct paradigms—“independent evolution” and “co-evolution”—the likelihood of the observed data under a given set of model parameters being tractable in the former and intractable in the latter. [Our derivation of the “independent evolution” likelihood function is given in the Appendix to this paper.]

3.1.1.1 The “Independent Evolution” Case As the name suggests in the “independent evolution” case we suppose that neither the birth nor morphological transformation history of any galaxy are ever coupled to those of another. Simulation from our stochastic model under this assumption for a given set of input parameters is then simply a matter of applying the above probabilistic transition rules to generate one-by-one a mock morphology at the observed redshift for each object in our benchmark sample as follows.

First, the birth time of the galaxy at hand (i.e., the epoch at which its stellar mass finally exceeds $10^{11} M_{\odot}$) is drawn from the interval $t_6(=0) \leq t_{\text{birth}} \leq t_{\text{obs}}$ according to the waiting time distribution dictated by its assumed (increasing-rate, non-homogeneous) Poissonian form (as derived in the Appendix to this paper). The birth class is then assigned as either Type III or Type IV, with the probability of the latter given by $\frac{W\lambda_m(t_{\text{birth}})}{\lambda_b(t_{\text{birth}})}$. To compute this ratio we must also sample a value for each of the nuisance parameters, K , γ and W , according to $f_{K,\gamma} \sim \mathcal{N}_{\text{Trunc.}}([-4.1, 0.65]', [0.06^2, 0.1^2]; \rho = 0.05; 0 < \gamma < 1)$ and $f_W \sim \mathcal{N}_{\text{Trunc.}}(\mu = 0.5, \sigma = 0.2; W > 0)$ respectively (where $\mathcal{N}_{\text{Trunc.}}$ represents the truncated Normal distribution). The number of mergers, n_{merge} , experienced between birth and observation is then drawn from the Poisson distribution specified by rate, $\Gamma_m^* = \int_{t_{\text{birth}}}^{t_{\text{obs}}} \frac{\lambda_m(t)}{\lambda_b(t)} dt$. If $n_{\text{merge}} \neq 0$ the corresponding epoch of last major merger is identified by sampling from the relevant waiting time distribution (also derived in the Appendix). The manifest duration of the resulting post-merger (Type IV) irregular state is then drawn directly from the Gamma distribution with scale, $1 + 100\tau_{\text{irr morph}}$, and shape coefficient, 100; and if encompassed within the remaining time until observation the galaxy is assigned either Type I or Type II morphology, with the probability of the latter set by $P_{\text{Sph+D remnant}}$ (otherwise it finishes the simulation as a Type IV). Finally, galaxies born as Type III disks and experiencing no major mergers may yet evolve to Type II via secular evolution, determined likewise by comparing an evolutionary period drawn from the Gamma distribution with scale, $1 + 50\tau_{\text{sec ev}}$, and shape coefficient, 50, against the time available between birth and observation.

Simulation from our model is thus inherently stochastic—i.e., the internal assignment of birth times, most recent merger times, and so on (and thereby the output assignment of final morphologies) will vary from run to run at fixed input. In the SMC approach to ABC (Chopin 2002; Del Moral et al. 2006; Sisson et al. 2007; Drovandi & Pettitt 2010) the effects of this stochasticity are accounted for in an efficient, consistent manner through

the iterative application of the key rejection and resampling/refreshment steps described in Section 3.2.

As noted earlier a characteristic feature of our model in the “independent evolution” case is that the likelihood, $P(\mathbf{y}|\boldsymbol{\theta})$, of the observed data under a given set of input parameters is, in fact, tractable (whereas in the “co-evolution” case it is not). For expository purposes this allows us to reconstruct via standard (“tractable likelihood”-based) Bayesian computational methods (namely, MCMC) the “true” posterior probability density of our model parameters (modulo the inherent variance of MCMC simulation) as a benchmark for comparison against our ABC results. The derivation of this likelihood function is rather involved so we present details separately in the Appendix (along with a description of the MCMC scheme employed). The resulting “true” posteriors are, however, presented here in Figure 4 for reference during our discussion of summary statistics in Section 3.3 and the accuracy of our ABC posteriors in Section 3.4.

3.1.1.2 The “Co-Evolution” Case Though it ensures the likelihood tractability required for our demonstration of the robustness of SMC ABC with respect to MCMC in Section 3.4 below our initial assumption that galaxy evolution proceeds independently across our entire sample may not be physically realistic. A number of recent studies probing high redshift clusters and proto-clusters out to $z \sim 1.5$ –3 (Doherty et al. 2010; Hatch et al. 2011; Papovich et al. 2011; Spitler et al. 2011; Tanaka et al. 2011) and beyond (Capak et al. 2011; Carilli et al. 2011) have presented evidence to suggest that the evolutionary histories of intermediate mass galaxies within these early over-densities are in fact highly correlated such that they consistently achieve peak star formation earlier than their counterparts of similar mass in the field. Indeed such early biasing by environment of star formation and mass accretion constitutes a fundamental prediction of hierarchical formation theory under Λ CDM (Springel et al. 2005; Overzier et al. 2009) and should already be manifest in the spatial distribution of the first generation of (re-)ionising sources (Kramer et al. 2006). As usual though, empirical results for galaxies at the top end of the stellar mass function remain limited due to the intrinsic rarity of these systems.

The beauty of ABC, of course, is that it permits the study of arbitrarily complex stochastic models irrespective of likelihood tractability, allowing one to relax such simplifying assumptions as that of “independent evolution” in the present example. In this Section we thus outline a “co-evolution” case of our model in which a physically plausible coupling is introduced into the formation times of galaxies in close pairs and small groups, and in Section 3.4 we explore the impact of this coupling on our ABC posteriors.

In Figure 5 we illustrate the nature of spatial clustering amongst the massive ($M_{\text{gal}} > 10^{11} M_{\odot}$) galaxies of our B11 (CANDELS/EGS) sample at $1.5 < z < 3$, representing their 3D distribution in observed right ascension (RA), declination (DEC), and redshift via 2D projections in comoving distance along the line of sight (LOS) and the axes of RA and DEC, alternately. Neighbouring systems separated by no more than 2.5 Mpc—a conservative linking scale for proto-group-sized over-densities in the early Universe (cf.

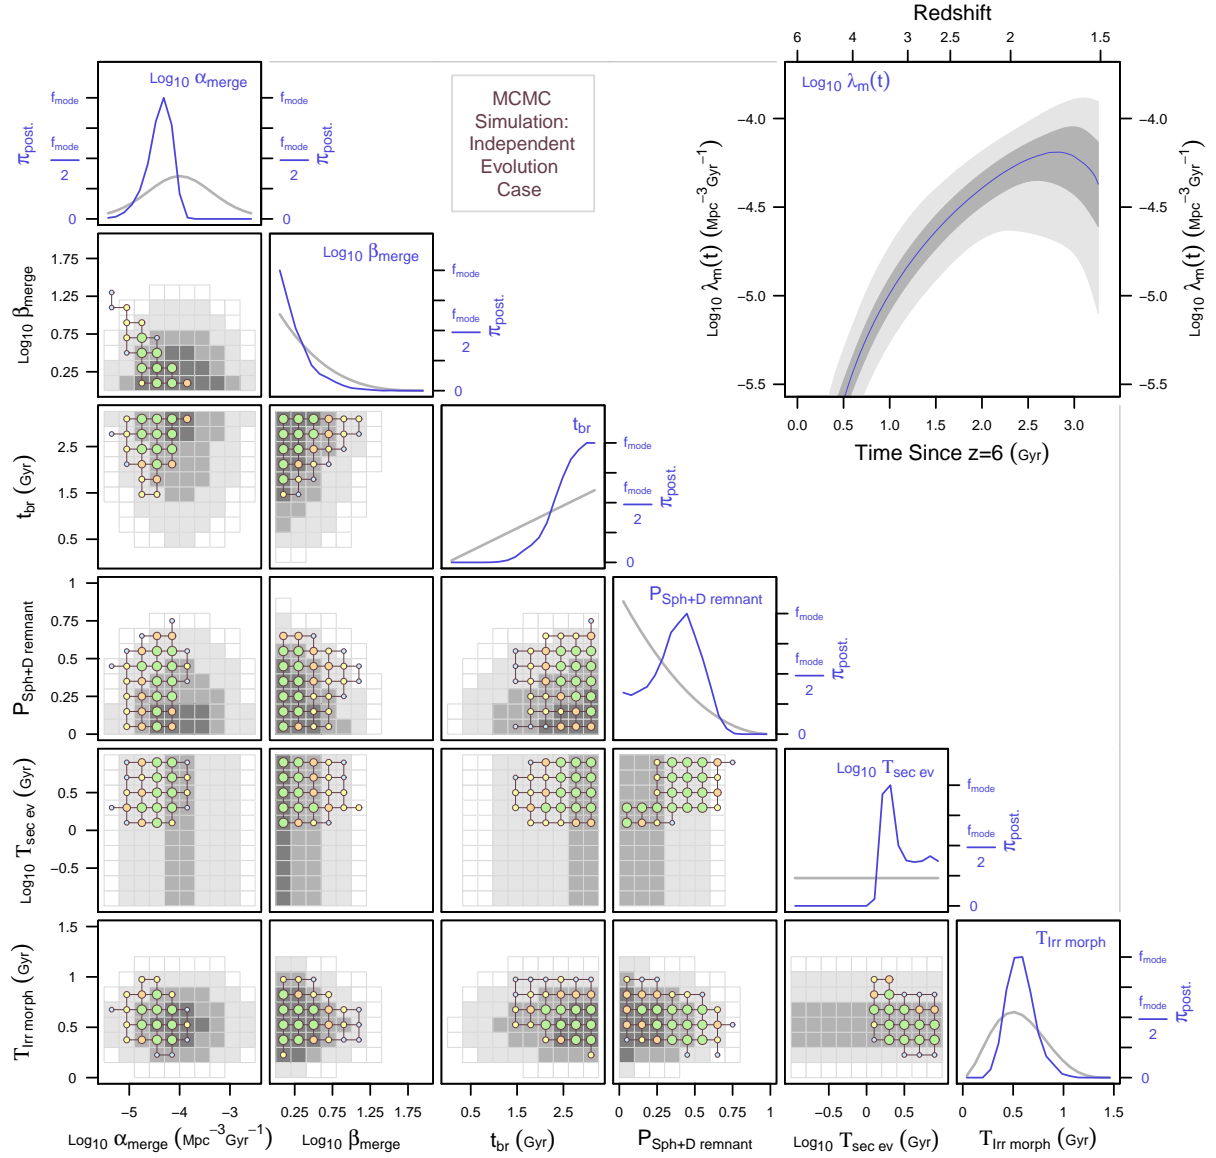


Figure 4. Benchmark posterior probability densities for the key parameters of our stochastic model of morphological transformation at high redshift (in the “independent evolution” case) recovered from “tractable likelihood”-based MCMC simulation. In each of the main diagonal panels we compare the marginal posterior density of a single parameter (in blue) against its prior (in grey), while in each of the off-diagonal panels below we extend this comparison to the joint density formed by pairing that parameter against one of its peers. For the latter visualisation we employ a lattice of variable-sized points to trace the MCMC posterior on a scale of 1, 2.5, 7.5, and 15 times some appropriate baseline probability density, while grey-shaded tiles map the corresponding prior on an identical scale. In the upper right panel we plot the (pointwise) 1σ and 3σ credible intervals and median curve (in dark grey, light grey, and blue respectively) for the $M_{\text{gal}} > 10^{11} M_{\odot}$ merger rate, $\lambda_m(t)$, deriving from our (joint, marginal) posterior densities on α_{merge} , β_{merge} , and t_{br} .

Capak et al. 2011 and references therein)—are marked accordingly and highlight the diversity of high redshift “environments” in the survey volume, with eleven simple pairings and one threesome identified at the adopted linking scale and a majority of relatively isolated systems. Perhaps the most striking feature on first inspection of this plot, however, is the apparent void at ~ 150 Mpc distance from the lower bound of our sample at $z \sim 1.5$. Examination of the B11 photometric redshifts for the full EGS, however, reveals no indication of this under-density extending across the wider field, reassuring us that it is indeed an imprint of cosmic variance (cf. Trenti & Stiavelli 2008; Driver & Robotham 2010)

within the small volume probed by the present dataset—and not the result of a systematic bias in the adopted SED fitting algorithm, for instance.

To explore the impact of small scale clustering on the “birth” time distribution as defined in our stochastic model for morphological transformation we refer to the publicly-available mocks from the De Lucia & Blaizot (2007) SAM embedded in a small volume of the Millennium Simulation¹¹ of comparable size to that probed by the B11 dataset.¹² In

¹¹ URL(<http://galaxy-catalogue.dur.ac.uk:8080/Millennium/>).

¹² Recall that for the purposes of the present analysis we treat




Figure 5. The spatial distribution of massive ($M_{\text{gal}} > 10^{11} M_{\odot}$) galaxies in our B11 (CANDELS/EGS) sample at $1.5 < z < 3$, projected in terms of comoving distance onto the LOS-RA and LOS-DEC planes. The blue circles overplotted highlight our eleven pairs and one threesome of galaxies neighbouring within 2.5 Mpc. The colour/symbol code for galaxy types is identical to that of Figure 1. For reference we also mark the LOS axis on a scale of time since $z = 6$ (the origin point of our model). [See the text for a comment on the apparent void at ~ 150 Mpc.]

line with our suspicion from Section 3.1 of an ~ 0.1 dex bias in the B11 stellar masses our first discovery here is that the number density of simulated galaxies in these mocks at a cut-off of $M_{\text{gal}} > 10^{11} M_{\odot}$ is far below that of our CANDELS/EGS sample, but may be brought into reasonable agreement if we revise our selection down to (at least) $M_{\text{gal}} > 10^{10.9} M_{\odot}$. Following this adjustment we identify 12 pairs, 2 threesomes, 3 foursomes, and even a five member configuration in this $\sim 7 \times 10^5 \text{ Mpc}^3$ “snapshot” of the De Lucia & Blaizot (2007) SAM at $z \sim 1.5$. Interestingly, whilst we do find a strong *correlation* between the birth times of galaxies in close associations—with a median absolute difference of only ~ 0.3 Gyr for neighbours within 2.5 Mpc compared against ~ 0.65 Gyr for randomly assigned pairs—despite some theoretical expectation there exists negligible evidence in these mocks for a systematic *bias* in this specific aspect of galaxy formation. Hence, we do not attempt to induce one arbitrarily into our stochastic simulations; instead we focus here simply on reproducing the observed correlation using the following modified sampling scheme.

Rather than drawing independent birth times one-by-one for each galaxy in our sample as in the “independent evolution” case described above, in the “co-evolution” case of our model we in fact generate a complete set of birth times at the start of the simulation and distribute these thereafter with an environmental dependence. We achieve this by drawing from the standard Uniform distribution a primary set of 126 points, one for each galaxy in the B11 CANDELS/EGS dataset, plus a secondary set of 12 points, one for each close pair or threesome earlier identified (see Figure 5)—the latter serving as “latent variables” for the coupling of birth times in these associations. To each galaxy in a close pair or threesome we then allocate a single point from the primary set with selection probability proportional

to the inverse square of distance between that point and the corresponding latent variable from the secondary set. The remaining points in the primary set are then randomly allocated with equal selection probability to the many isolated galaxies of our dataset. Each point on the interval $[0,1]$ thus assigned is then transformed to a birth time for its matching galaxy through multiplication by the relevant $\Lambda_b(t_{\text{obs}})$ (see, for reference, the Appendix to this paper) followed by inversion to recover t_{birth} . We have confirmed by repeat simulation that this process of birth time sampling consistently produces median absolute birth time differences for both neighbouring galaxies and random pairings of similar magnitude to those of the De Lucia & Blaizot (2007) SAM described above. Note also that by allocating from an initial uniform sample in this manner we naturally preserve the mean build up rate of $\Phi_{>10^{11} M_{\odot}}(z)$ specified by our fit of K and γ in $\lambda_b(t)$ against the available observational data from K09, M10, BR11, C11, M11, and C12 (see the lefthand panel of Figure 2).

3.2 SMC ABC: Sequential Monte Carlo Approximate Bayesian Computation

As mentioned in the Introduction, Approximate Bayesian Computation (cf. Tavaré et al. 1997; Pritchard et al. 1999; Beaumont et al. 2002; Sisson et al. 2007; Wilkinson 2008; Csilléry et al. 2010; Drovandi & Pettitt 2010, 2011) offers a rigorous statistical framework for estimating the posterior probability densities of key scientific parameters under complex models for which the likelihood of the observed data may be entirely intractable (thus prohibiting application of the standard MCMC approach, for example). To conduct ABC one requires only a stochastic model from which the observed data, \mathbf{y} , are believed to be a random draw given some unknown set of intrinsic (true) input parameters, and a discrepancy measure for the comparison of simulated data, \mathbf{y}_s , against observed, $\rho(S(\mathbf{y}), S(\mathbf{y}_s))$, typically based on a set of low order summary statistics, $S(\cdot)$. Following Drovandi & Pettitt (2010) the aim of ABC may be stated formally as the recovery of unbiased samples from the distribution described by the approximate joint posterior density,

$$f(\boldsymbol{\theta}, \mathbf{y}_s | \rho(S(\mathbf{y}), S(\mathbf{y}_s)) \leq \epsilon_T) \propto \frac{f(\mathbf{y}_s | \boldsymbol{\theta}) \pi(\boldsymbol{\theta})}{\times \mathbf{1}_{\rho(S(\mathbf{y}), S(\mathbf{y}_s)) \leq \epsilon_T}}, \quad (1)$$

where $\boldsymbol{\theta}$ represents a vector of unknown model parameters, $\pi(\boldsymbol{\theta})$ the prior density on those parameters, and ϵ_T some target tolerance for the specified discrepancy measure between simulated and observed data. The indicator function $\mathbf{1}_{\rho(S(\mathbf{y}), S(\mathbf{y}_s)) \leq \epsilon_T}$ assumes value unity for simulated–observed dataset pairs with metric distance below this tolerance, and zero otherwise.

The archetypal scheme for random sampling from the distribution defined by Equation 1 is that of rejection ABC (cf. Pritchard et al. 1999) in which one draws a large sample of N trial parameter vectors, $\boldsymbol{\theta}_i$ ($i = 1, \dots, N$), from the prior, $\pi(\boldsymbol{\theta})$, simulates a corresponding dataset for each, \mathbf{y}_s , and then rejects all $\boldsymbol{\theta}_i$ for which the discrepancy between simulated and observed data exceeds some target tolerance, i.e., $\rho(S(\mathbf{y}), S(\mathbf{y}_s)) > \epsilon_T$. Unfortunately, except in (unrealistically) fortuitous (or trivial) circumstances in which the prior is already very close to the posterior, when exploring high dimensional parameter spaces ($N_{\text{par}} \gtrsim 3$) under

the time of “birth” in our model as the epoch at which a galaxy first reaches the top end of the high redshift stellar mass function, whether via star formation or merging. To identify this point in the De Lucia & Blaizot (2007) mocks one must follow back the linked progenitor tree accordingly via SQL query.

rejection ABC one will often face an undesirable trade-off between an impractically low acceptance rate or an uncomfortably large tolerance—either way the approximation of the true posterior by the ABC posterior will likely be poor.

As a more efficient alternative the Sequential Monte Carlo (cf. Liu 2001; Chopin 2002) approach to ABC (Del Moral et al. 2006; Sisson et al. 2007; Drovandi & Pettitt 2010) proposes instead to simulate from Equation 1 step-wise by evolving a dynamic population of “particles” (with each particle representing a single vector of input model parameters) through a sequence of intermediate distributions characterised by $f(\mathbf{y}_s|\boldsymbol{\theta})\pi(\boldsymbol{\theta})\mathbf{1}_{\rho(S(\mathbf{y}), S(\mathbf{y}_s)) \leq \epsilon_t}$ for $t = 1, \dots, T$, indexing a series of non-increasing targets, ϵ_t . The two key stages of the SMC algorithm are thus: (i) rejection of the most discrepant particles under each target; followed by (ii) resampling from amongst the least discrepant particles with some refreshment mechanism applied to maintain particle diversity. SMC may therefore also be referred to as “particle filtering” or “population Monte Carlo” (PMC; for examples of the recent adoption of PMC [SMC] in a cosmological context see Wraith et al. 2009, Kilbinger et al. 2010, and references therein).

In this study we employ for rejection the self-generating target strategy of Drovandi & Pettitt (2010) in which the sequence of incremental targets, ϵ_t , is chosen on run-time such that a fixed fraction, α , of all particles are dropped at each iteration (herein $\alpha = 0.75$). We then restore the particle population to its full operating size, N , by resampling with replacement from amongst the remaining $(1 - \alpha)N$ particles. Population refreshment is achieved by application of an MCMC kernel to these replicates. As in Drovandi & Pettitt (2010) we favour the use of a self-refining, Metropolis-Hastings proposal distribution based on the current particle sample mean vector, $\bar{\boldsymbol{\mu}}$, and covariance matrix, $\bar{\boldsymbol{\Sigma}}$. For the specific model at hand we adopt a truncated, multivariate T distribution of degree 10 with the truncation bounds set by the support of our prior densities (as detailed in Section 3.1). At each iteration of the SMC algorithm this MCMC kernel is run a fixed number of times, R , to (hopefully) produce genuine refreshment in a fraction, c , of the resampled particles (with some particles, of course, likely to be moved multiple times). The requisite R is here estimated according to the empirical efficiency, p_{acc} , of the previous MCMC kernel as $R = \frac{\log(1-c)}{\log(1-p_{\text{acc}})}$. Note that the “likelihood ratio” in the corresponding MCMC acceptance computation is simply $\mathbf{1}_{\rho(S(\mathbf{y}), S(\mathbf{y}_s(\boldsymbol{\theta}_{\text{proposed}}))) \leq \epsilon_t}$, i.e., all trial particles for which the simulation produces a mock dataset with discrepancy distance falling below the current target are assured a non-zero probability of acceptance.¹³ Our final target, ϵ_T , is defined pragmatically (with respect to the limitations of our computational resources) as that ϵ_t for which a further SMC rejection-resampling step would incur more than $R_{\text{max}} \sim 100$ applications of this MCMC kernel.

Treatment of Nuisance Parameters A particular feature of the model adopted in this study is the appear-

ance of three nuisance parameters—essential inputs known only to a limited accuracy from previous studies, but unable to be robustly constrained via the present observational dataset. Namely, $\boldsymbol{\Theta} = \{K, \gamma, W\}$ where $f_{K,\gamma} \sim \mathcal{N}_{\text{Trunc}}([-4.1, 0.65]', [0.06^2, 0.1^2]; \rho = 0.05; 0 < \gamma < 1)$ and $f_W \sim \mathcal{N}_{\text{Trunc}}(\mu = 0.5, \sigma = 0.2; W > 0)$. Upon each simulation from our model for a single particle, $\boldsymbol{\theta}_i$, we draw a random value for each of K , γ , and W from their respective distributions for use in that one instance. Heuristically this amounts to approximating the integral, $\int_{\Omega(\boldsymbol{\Theta})} f(\mathbf{y}_s|\boldsymbol{\theta}_i, \boldsymbol{\Theta})f(\boldsymbol{\Theta})d\boldsymbol{\Theta}$, by a single Monte Carlo sample, yet produces through the power of the SMC ABC algorithm (i.e., via inference over a particle population *en masse*) a reasonable approximation (converging asymptotically) to sampling from Equation 1.

3.3 Refinement of Summary Statistics

As explored in a number of recent papers careful selection of the summary statistic(s) used to evaluate the discrepancy between simulated and observed data in ABC is of paramount importance to achieving accuracy and efficiency with this algorithm, whether employed for the purpose of parameter estimation (e.g. Joyce & Marjoram 2008; Nunes & Balding 2010; Fearnhead & Prangle 2012) or, more challengingly¹⁴, Bayesian model choice (Robert et al. 2011; Marin et al. 2011; Barnes et al. 2011). Indeed the same is true for all approaches to inference from complex models for which sufficient statistics are unavailable, including those of “approximate likelihood” (Wood 2010) and model emulation (cf. Bower et al. 2010 and references therein), and should thus not be thought of as a unique concern for ABC-based analyses.

For some stochastic models (including that of the present case study in morphological transformation; cf. Section 3.1) the nature of the output data may well be such that only a few basic modes of summary are of any likely value, while for others (e.g. the Ricker map case study of Wood 2010; and most SAM-based studies of galaxy formation¹⁵) there may in fact be many in a rich hierarchy of complexities—either way though it will rarely be profitable to carry more than one or two into a full ABC analysis, owing to the direct relationship between Monte Carlo error and summary dimension (cf. Beaumont et al. 2002; Fearnhead & Prangle 2012). Hence in this Section

¹⁴ Application of ABC to the problem of Bayesian model choice (cf. Grelaud et al. 2009; Toni & Stumpf 2010) is far from straightforward as an unfortunate summary statistic selection can lead to disastrously incorrect Bayes factors, even asymptotically (Robert et al. 2011). Recently though, Marin et al. (2011) have made substantial progress in this field by establishing necessary and sufficient conditions on the validity of candidate summary statistics for this purpose.

¹⁵ With the choice of summary statistic typically (re-)cast in terms of the choice of reference dataset in these astronomical studies—namely, whether to constrain, for instance, against the luminosity function (Bower et al. 2010; Lu et al. 2011; Cirasuolo et al. 2010), the Tully-Fisher relation (van den Bosch 2000; Tonini et al. 2011), the mass-metallicity relation (Pipino et al. 2009), and/or the black hole–bulge mass relation (Henriques et al. 2009).

¹³ When using a non-symmetric proposal distribution as in the present case the ratio of sampling densities joins, of course, the likelihood ratio and the prior density ratio in computing the full MCMC acceptance probability.

we complete our exposition of the ABC technique with a demonstration of two contemporary approaches to this particular selection problem: first, we apply Nunes & Balding’s (2010) two-stage procedure of distributional entropy and MRSSE¹⁶ minimisation to identify the optimal choice from amongst a candidate set of four “naïve” summary statistics for our morphological dataset; and second, we employ the so-called “semi-automatic” scheme of Fearnhead & Prangle (2012) to build an alternative set of summary statistics optimised with respect to the recovery of posterior means, validating their performance in comparison against the former.

Minimum Entropy/MRSSE-Based Selection The most natural mode of summary for population demographic data in the context of extra-galactic astronomy is, of course, by way of type counts or type fractions in similar-sized bins of redshift (see Oesch et al. 2010 and Buitrago et al. 2011 for recent examples). Important considerations when compiling such summary data for the purpose of ABC analysis are then the number and placement of bins to use and the weights one should assign to the type counts/fractions observed therein. As a “naïve” first attempt at constraining the parameter space of our model in the “independent evolution” case we thus trial four alternative summary statistics based on progressively finer subdivisions of the B11 CANDELS/EGS dataset by redshift. Mathematically, we define this class of “naïve” summary statistics, \mathbf{S} , via the generic column vector

$$\mathbf{S} \equiv S(\mathbf{y}) = \{f_{\text{I}}^{(1)}, f_{\text{II}}^{(1)}, f_{\text{III}}^{(1)}, f_{\text{IV}}^{(1)}, \dots, f_{\text{I}}^{(k)}, f_{\text{II}}^{(k)}, f_{\text{III}}^{(k)}, f_{\text{IV}}^{(k)}\}',$$

with $f_{\text{I}}^{(1)}$ the fraction of Type I galaxies in the first of k redshift bins subdividing equally the interval $1.5 < z < 3$, $f_{\text{II}}^{(1)}$ the fraction of Type II galaxies in the aforementioned bin, and so on. Adopting equal significance weights across all bins, we thus establish a complete discrepancy distance,

$$\rho(S(\mathbf{y}), S(\mathbf{y}_s)) = \sqrt{(\mathbf{S}_{\text{obs}} - \mathbf{S}_{\text{sim}})'(\mathbf{S}_{\text{obs}} - \mathbf{S}_{\text{sim}})}. \quad (2)$$

Following the two-stage procedure of Nunes & Balding (2010) we begin the evaluation of our four “naïve” candidates, $\mathbf{S} : k = \{1, 3, 6, 12\}$, by computing the fourth-nearest neighbour entropy¹⁷ of the posterior distribution resulting from simple rejection ABC under each—the goal here being to exploit the (approximate) inverse relationship between entropy and information in order to identify the most “informative” summary statistic with regard to inference of the model parameters at hand. In the present round of rejection ABC experiments we accept only the 100 least discrepant particles of an initial sample of 5,000 drawn from the prior, and we compute the associated entropy statistic for each according to the formula,

¹⁶ MRSSE: Mean square Root Sum of Standard Errors.

¹⁷ One may note an intriguing similarity between the use of fourth (or fifth) nearest neighbour-based estimators in both statistical studies of distributional entropy and in astronomical studies of large-scale environment (cf. Baldry et al. 2006)—though it is unlikely there exists an underlying significance to this beyond the desirable error properties of the $n \sim 4$ -5 choice (Singh et al. 2003).




Figure 6. Evaluation of candidate summary statistics for model-data comparison across a range of binning schemes ($k = 1, 3, 6, 12$) for both our “naïve” \mathbf{S} and optimised \mathbf{S}^* (the latter explained later in this Section) via the twin diagnostics of (Left:) distributional entropy and (Right:) (M)RSSE (cf. Nunes & Balding 2010). In each instance the marked datapoint reveals the median, and the error bars a corresponding 95% confidence interval, recovered from six rounds of rejection ABC (i.e., selection of the 100 least discrepant particles out of an initial 5,000 drawn from the prior density). Note that as the posterior means of our model parameters under the “independent evolution” case studied here have already been well approximated via our earlier (“tractable likelihood”-based) MCMC simulation we have employed these directly to estimate a pseudo-RSSE, rather than forming an MRSSE from repeated ABC runs against simulated datasets “close” to the real one as in the canonical Nunes & Balding (2010) procedure.

$$\hat{H} = \log \left[\frac{\pi^{N_{\text{par}}/2}}{\Gamma(N_{\text{par}}/2 + 1)} \right] - \psi(4) + \log n + \frac{N_{\text{par}}}{n} \sum_{i=1}^n \log R_{i,4}$$

(Singh et al. 2003; with $N_{\text{par}} [= 6]$ representing the dimension of our model parameter space, $n [= 100]$ the number of accepted particles, $\Gamma(\cdot)$ and $\psi(\cdot)$ the “gamma” and “digamma” functions, respectively, and $R_{i,4}$ the fourth nearest neighbour distance). There exists, of course, a certain subjectivity in the choice of scaling for each parameter in the computation of $R_{i,4}$; one option would be to first standardise all parameters on the interval $[0, 1]$, however, in this case we prefer instead to standardise against the diagonal matrix of our prior variances, \mathbf{V} , such that $\|\boldsymbol{\theta}_i, \boldsymbol{\theta}_j\| = \sqrt{\boldsymbol{\theta}_i' \mathbf{V}^{-1} \boldsymbol{\theta}_j}$. By repeating the rejection ABC process six times for each k one may estimate both the median entropy and matching 95% confidence interval (from the range) under that particular binning scheme. The results of this analysis are presented in the lefthand panel of Figure 6.

Interestingly, although one might, at face value, expect a monotonic relationship of decreasing posterior distributional entropy with increasing k on the basis that finer binning should break any false degeneracies in the posteriors recovered from ABC runs with fewer bins (i.e., in some sense increase the information return), this is not necessarily true in practice owing to the simultaneous increase in Monte Carlo error, as the present example demonstrates. Although we do observe a slight decrease in distributional entropy upon moving from one to three redshift bins the opposite is true for our six and twelve bin trials under this particular mode of summary (see the lefthand panel of Figure 6).

Following identification of the entropy-minimising \mathbf{S} (here $k = 3$) Nunes & Balding (2010) recommend a second (rather computationally expensive) round of evaluation against the formal optimality criterion of Mean square Root Sum of Squared Errors (MRSSE), to be estimated via a new




Figure 7. Comparison of the marginal posterior density for α_{merge} in the “independent evolution” case of our model recovered from SMC ABC under two of our “naïve” summary statistics, $\mathbf{S} : k = 3$ and $\mathbf{S} : k = 12$). In each case the dotted line represents the posterior after one rejection-resampling iteration ($m = 1$) of the SMC ABC algorithm, the solid datapoints the same after four iterations ($m = 4$), and the solid blue line the benchmark posterior from “tractable likelihood”-based MCMC. At the bottom of each panel we indicate also the position of the posterior mean for this parameter under each of our SMC ABC runs as well as from our MCMC benchmark (the latter highlighted with a blue diamond).

series of rejection ABC analyses against simulated datasets constructed under parameter vectors revealed by the original ABC runs for the entropy-minimising \mathbf{S} as (likely to be) “close” to those responsible for the observed dataset. Since the (marginal) posterior mean of each model parameter under the “independent evolution” case studied here has already been well approximated via our earlier (“tractable likelihood”-based) MCMC simulation (see Section 3.1) we may take a shortcut to the truth (and vastly reduce our computational burden) by employing these directly to estimate alternative pseudo-RSSE scores for each of our previous rejection ABC runs. The results of this analysis are presented in the righthand panel of Figure 6. The relationship between pseudo-RSSE score and binning k observed here mirrors closely that exposed by our original entropy evaluation, validating $\mathbf{S} : k = 3$ as the optimal choice from amongst our candidate set of “naïve” summary statistics.

The value of the above optimisation procedure may easily be appreciated from inspection of Figure 7 in which we compare the marginal posterior for one of our key model parameters, α_{merge} , recovered from full SMC ABC analysis under \mathbf{S} with $k = 3$ against that for $k = 12$. As in all SMC ABC runs reported herein we use a population of 10,000 particles, iterated through an adaptive threshold defined by a rejection rate of $\alpha = 0.75$, followed by MCMC kernel-based resampling with a goal refreshment rate of no less than $c = 0.90$. A total of four iterations were achieved under this scheme before the limits of our computational resources were reached (at $R > R_{\text{max}} [= 100]$ required applications of the MCMC kernel for such a level of refreshment). Although both (marginal) posteriors appear rather similar after only one iteration it is evident by the fourth (and final) iteration that the $\mathbf{S} : k = 3$ statistic has produced a particle population tracing far more faithfully the density of our benchmark (“tractable likelihood”-based) MCMC simulation than its $\mathbf{S} : k = 12$ counterpart.

The “Semi-Automatic” Scheme In an important contribution to the ABC literature Fearnhead & Prangle (2012) have recently demonstrated that the optimal sum-

mary statistic for estimation of model parameters under quadratic loss (i.e., optimality with respect to the recovery of posterior means) is simply the conditional expectation function, $E(\theta|\mathbf{y})$. As a direct consequence the authors were thereby able to propose and justify a regression-based algorithm for the direct construction of well-behaved summary statistics, allowing one (in principle) to by-pass the above process of searching through a “naïve” set of often unsatisfactory candidates. To implement their so-called “semi-automatic” scheme one must first generate a “reasonably large” sample of model parameter-simulated dataset pairs spanning a “relevant volume” of parameter space. In lower dimensional analyses one may simply draw this sample directly from the prior, though the posterior density from a trial ABC run with some “naïve” summary statistic will generally serve as a superior starting point. Least-squares-based fitting to this reference dataset of the relation, $\theta_i = \beta_0^{(i)} + \beta_1^{(i)} f(\mathbf{y}) + e_i$, for each model parameter, $\theta_i \in \Theta$, yields the optimal summary statistic, $\mathbf{S}^* = \beta_0 + \beta_1 f(\mathbf{y})$.¹⁸ Here e_i denotes a symmetric error term of zero mean and $f(\mathbf{y})$ some vector-valued function of the data, which for $\mathbf{y} \in \mathbb{R}^p$ will be typically of the form $f(\mathbf{y}) = \mathbf{y}$, $f(\mathbf{y}) = \{\mathbf{y}, \mathbf{y}^2\}'$, or similar—the lack of a universally appropriate algorithm for defining this regression function being the chief cause for classification of the above scheme as “semi-” rather “fully” automatic.

Since the raw output, \mathbf{y} , from our stochastic model for high redshift morphological transformation is, in fact, multinomial (rather than real-valued) we adopt here (for the purposes of computational efficiency) a modified regression function of form, $f(\mathbf{y}) = S(\mathbf{y})$, with $S(\cdot)$ denoting as above the compilation of type fractions in fixed bins of redshift. Under this adaptation of the Fearnhead & Prangle (2012) approach the magnitude of each component in each fitted $\beta_1^{(i)}$ may be considered a weight for the importance of that type fraction and redshift bin in estimating the corresponding (i -th) model parameter. As in our earlier computation of fourth-nearest neighbour distances we employ our prior variance matrix, \mathbf{V} , to establish the full discrepancy measure under this new summary statistic¹⁹,

$$\rho(S^*(\mathbf{y}), S^*(\mathbf{y}_s)) = (\mathbf{S}_{\text{obs}}^* - \mathbf{S}_{\text{sim}}^*)' \mathbf{V}^{-1} (\mathbf{S}_{\text{obs}}^* - \mathbf{S}_{\text{sim}}^*).$$

Fitting of β_0 and β_1 was achieved here by application of the `glm` and `step` routines in R to the 100 least discrepant particles from each of the six rejection ABC runs conducted earlier under our “naïve” summary statistic for $k = 3$ (resulting in a full calibration sample of 600 model parameter-simulated dataset pairs). Notably, the `step` routine in R makes use of the AIC (Akaike Information Criterion) statistic to restrict each fit to only those elements of \mathbf{S} contributing significantly to the prediction of θ_i .

Although $k = 3$ proved to be the optimal binning scheme for the set of “naïve” summary statistics examined above one cannot simply assume this to hold for the new \mathbf{S}^* , so we again examine the merits of each alternative,

¹⁸ In fact, as noted by Fearnhead & Prangle (2012), since in ABC analysis we are only interested in the difference, $\mathbf{S}_{\text{obs}}^* - \mathbf{S}_{\text{sim}}^*$, the vector, β_0 , may well be omitted from this above definition.

¹⁹ Another (well-motivated) alternative choice of scaling here would be the sample covariance matrix of the posterior particle population from our earlier run of SMC ABC under $\mathbf{S} : k = 3$.

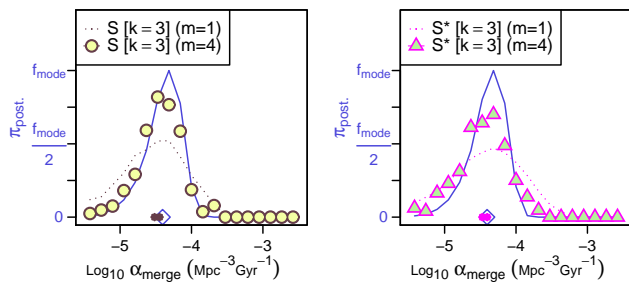


Figure 8. Comparison of the marginal posterior density for α_{merge} in the “independent evolution” case of our model recovered from SMC ABC under the alternative summary statistics, “naïve” $\mathcal{S} : k = 3$ and optimised $\mathcal{S}^* : k = 3$. In each case the dotted line represents the posterior after one rejection–resampling iteration ($m = 1$) of the SMC ABC algorithm, the solid data-points the same after four iterations ($m = 4$), and the solid blue line the benchmark posterior from “tractable likelihood”-based MCMC. At the bottom of each panel we indicate also the position of the posterior mean for this parameter under each of our SMC ABC runs as well as from our MCMC benchmark (the latter highlighted with a blue diamond).

$k = \{1, 3, 6, 12\}$, following Nunes & Balding (2010). The results of this analysis are overlaid against our measurements for the “naïve” \mathcal{S} in Figure 6. Reassuringly, with the exception of the most limited ($k = 1$) binning scheme these new summary statistics significantly out-perform the old in the pseudo-RSSE criterion for which they are designed. Despite our caution $k = 3$ does again appear to constitute the best choice of binning, though $k = 6$ is not far behind in accuracy and may also offer slightly lower entropy. In Figure 8 we compare the marginal posterior density recovered for the model parameter, α_{merge} , following a full run of SMC ABC under $\mathcal{S}^* : k = 3$ against that obtained earlier under $\mathcal{S} : k = 3$. Interestingly, though our “naïve” summary provides a visually “closer” fit to the shape (especially the width, i.e., standard deviation) of the benchmark MCMC density for this parameter, the optimised statistic does outperform it with regard to the recovery of the posterior mean (which it manages within a small tolerance after only a single iteration of the SMC ABC algorithm). Hence, given both its ease of implementation and its demonstrated effectiveness in the present analysis we can confidently recommend the “semi-automatic” scheme of Fearnhead & Prangle (2012) for summary statistic refinement.

A Note on One Alternative As mentioned in the introduction to this Section the subject of summary statistic selection for ABC analysis remains an active area of research in the statistical literature, hence it is worth reviewing here briefly another popular alternative we have neglected to demonstrate above for the sake of brevity. Namely, the “approximate sufficiency” algorithm of Joyce & Marjoram (2008) for iteratively building a master summary statistic from the union of randomly trialled candidates, with each new addition accepted only if it offers an improvement in parameter inference exceeding some threshold. This algorithm may be of particular interest for SAM-based studies of galaxy formation given the wide variety of available observational benchmarks from which summary statistics may be composed (see Footnote 15), though it has been criticised for a dependence on the (random) order in which the candi-

date statistics are tested at each application (i.e., the stated search procedure is far from exhaustive). Note also that although our above demonstration of the Nunes & Balding (2010) procedure is presented in terms of selecting a unique summary statistic from four evidently degenerate choices (i.e., $k = \{1, 3, 6, 12\}$) an optimal union of summary statistics may also be identified via the minimum entropy/MRSSE criterion, though possibly at great computational expense if the original set of basis candidates is large and there are many permutations of interest.

3.4 SMC ABC Posteriors for Our Stochastic Model of Morphological Transformation

In Figure 9 we present posterior probability densities for the key parameters of our stochastic model of morphological transformation at high redshift in the “independent evolution” case, as recovered from SMC ABC using our optimised summary statistic (cf. Section 3.3 above), $\mathcal{S}^* : k = 3$. The approximate solution shown here represents the state of a 10,000 particle population progressed through four rejection–resampling iterations with an $\alpha = 0.75$ rejection rate and a $c = 0.90$ target refreshment rate (cf. Section 3.2). Comparison against our “tractable likelihood”-based MCMC benchmark (for this tractable case of our model) presented earlier in Figure 4 highlights the robustness of the ABC approach. That is, without reference to the explicit likelihood function of the system at hand this simple procedure of strategic simulation and discrepancy thresholding has nevertheless produced a most satisfactory approximation to the true posterior, capturing the key features of each marginal and bivariate joint density under investigation. As is typical of ABC analysis though (e.g. Csilléry et al. 2010) the credible intervals thus derived (see, in particular, those for the $M_{\text{gal}} > 10^{11} M_{\odot}$ merger rate, $\lambda_m(t)$, shown in the top right panel) are noticeably broader than those of the benchmark “tractable likelihood”-based solution, owing to the two key sources of (“extra-Monte Carlo”) approximation error at play here: (I) the non-zero tolerance for the simulated–observed dataset discrepancy required to achieve a workable acceptance rate; and (II) the inherent degeneracy of summary statistic matching over full dataset matching (cf. Fearnhead & Prangle 2012, for instance).

In Figure 10 we present for comparison the SMC ABC posteriors recovered from the “co-evolution” case of our model for which the likelihood function is thoroughly intractable, and thus the standard toolbox of (“tractable likelihood”-based) MCMC simulation entirely is inaccessible. As described in Section 3.1 the only difference between our “co-evolution” and “independent evolution” simulations is that the birth times of galaxies in close associations are coupled in a manner believed “realistic” from analysis of the De Lucia & Blaizot (2007) SAM in the former but not in the latter. Hence it is perhaps unsurprising that the posteriors for this example are only subtly different from those presented above, with a slight decrease in confidence regarding the true value of the merger rate (i.e., the joint, marginal density of α_{merge} , β_{merge} , and t_{br}) and a slight increase in confidence regarding the merger visibility timescale (from $\tau_{\text{irr morph}} \approx 0.63 \pm_{0.18}^{0.20} [\sigma]$ to $\tau_{\text{irr morph}} \approx 0.53 \pm_{0.12}^{0.13}$). Nevertheless this demonstrated ability of ABC to handle models with intractable likelihoods, and thus to permit the deriva-

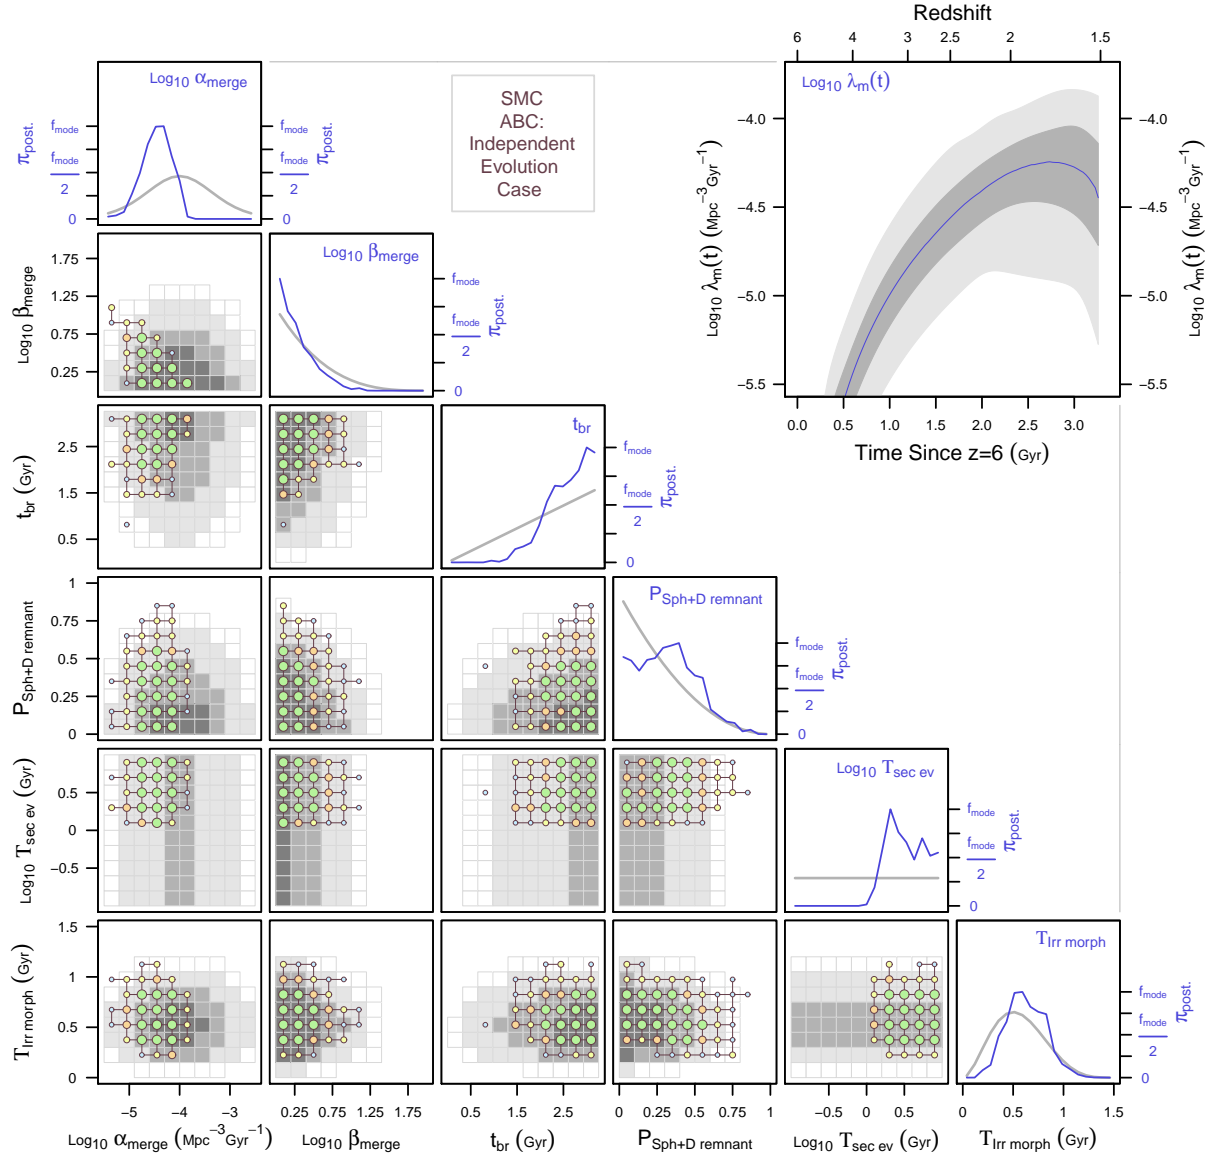


Figure 9. SMC ABC posterior probability densities for the key parameters of our stochastic model of morphological transformation at high redshift (in the “independent evolution” case). In each of the main diagonal panels we compare the marginal posterior density of a single parameter (in blue) against its prior (in grey), while in each of the off-diagonal panels below we extend this comparison to the joint density formed by pairing that parameter against one of its peers. For the latter visualisation we employ a lattice of variable-sized points to trace the SMC ABC posterior on a scale of 1, 2.5, 7.5, and 15 times some appropriate baseline probability density, while grey-shaded tiles map the corresponding prior on an identical scale. In the upper right panel we plot the (pointwise) 1σ and 3σ credible intervals and median curve (in dark grey, light grey, and blue respectively) for the $M_{\text{gal}} > 10^{11} M_{\odot}$ merger rate, $\lambda_m(t)$, deriving from our (joint, marginal) posterior densities on α_{merge} , β_{merge} , and t_{br} .

tion of robust Bayesian constraints from arbitrarily “realistic” (i.e., complex) simulations, offers a wealth of possibilities for astronomical studies far beyond the present example which cannot be over-stated.

4 ASTROPHYSICAL RESULTS & DISCUSSION

Having completed our exposition of the ABC algorithm in Section 3 above we take the opportunity here to explore a number of interesting astrophysical results arising from our

chosen case study in morphological transformation at high redshift. In Section 4.1 we compare our SMC ABC-based constraints on the evolving merger rate in the early Universe against recent estimates from the literature based on simple close pair and asymmetric galaxy counts, highlighting the superior informative power of the former over the latter. Then in Section 4.2 we discuss our (posterior) preference for merging over secular evolution as the dominant pathway to early bulge formation in the context of contemporary hydrodynamical and “semi-empirical” simulations.

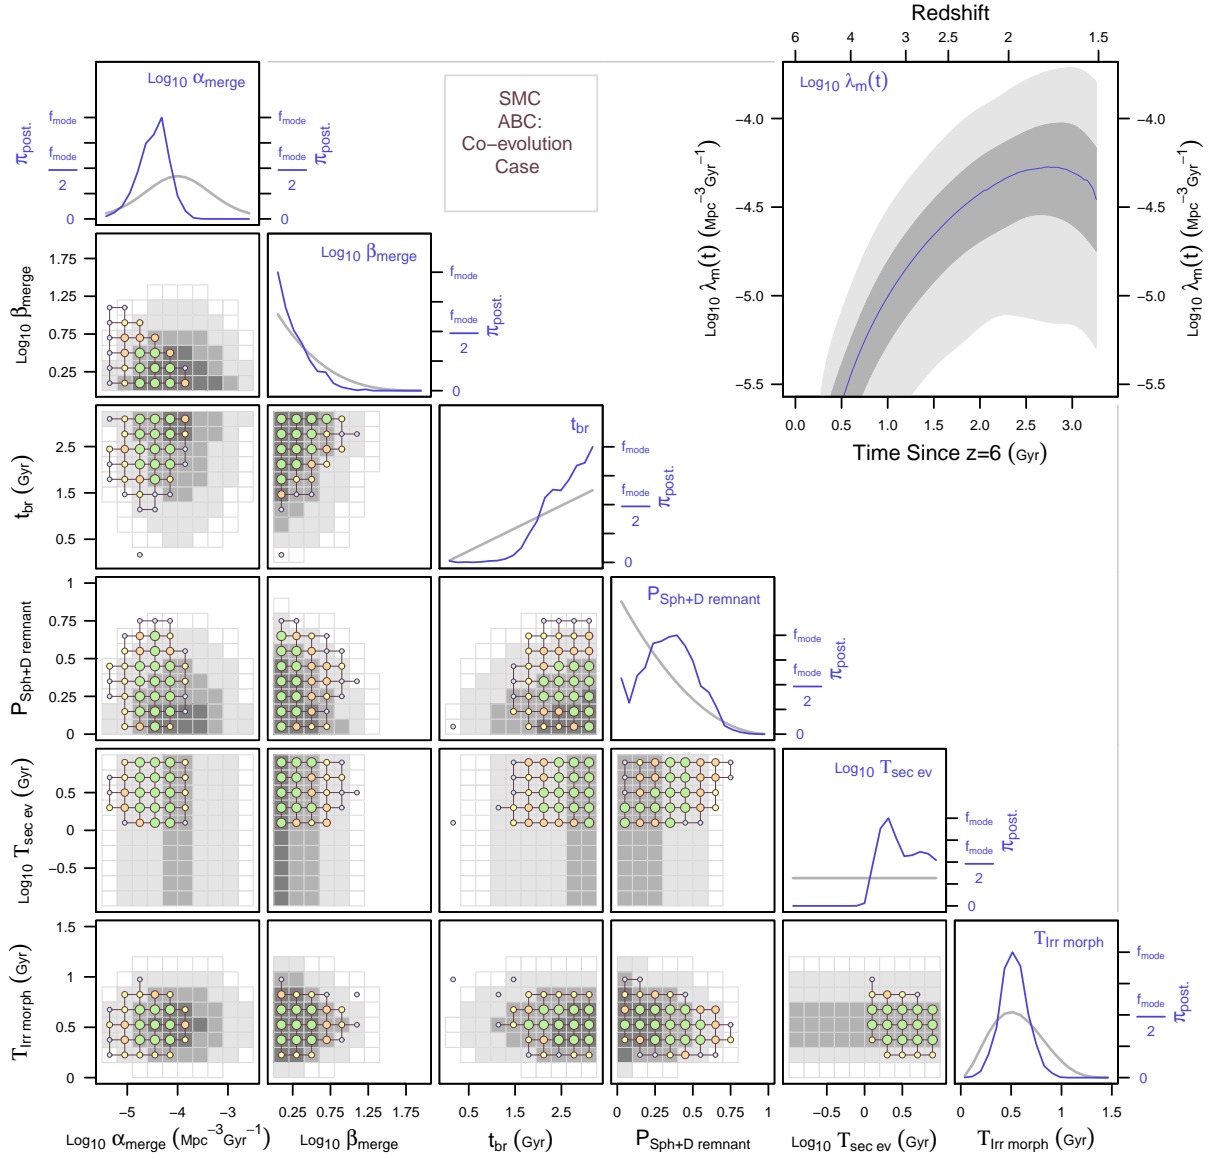


Figure 10. SMC ABC posterior probability densities for the key parameters of our stochastic model of morphological transformation at high redshift (in the “co-evolution” case). In each of the main diagonal panels we compare the marginal posterior density of a single parameter (in blue) against its prior (in grey), while in each of the off-diagonal panels below we extend this comparison to the joint density formed by pairing that parameter against one of its peers. For the latter visualisation we employ a lattice of variable-sized points to trace the SMC ABC posterior on a scale of 1, 2.5, 7.5, and 15 times some appropriate baseline probability density, while grey-shaded tiles map the corresponding prior on an identical scale. In the upper right panel we plot the (pointwise) 1σ and 3σ credible intervals and median curve (in dark grey, light grey, and blue respectively) for the $M_{\text{gal}} > 10^{11} M_{\odot}$ merger rate, $\lambda_m(t)$, deriving from our (joint, marginal) posterior densities on α_{merge} , β_{merge} , and t_{br} .

4.1 The Evolving Merger Rate at the Highest Redshifts

As mentioned in the Introduction to this paper the recent installation of WFC3 on HST has at last made accessible at high resolution the rest-frame optical morphologies of massive galaxies at the epoch of peak cosmic star formation and AGN activity ($z \sim 2$; Lilly et al. 1996; Madau et al. 1996; Oesch et al. 2012; Warren et al. 1994), opening a unique window into the structural assembly of this first generation of Hubble sequence analogues (Cameron et al. 2011b; Conselice et al. 2011a; Szomoru et al. 2011a). A particu-

lar motivation for our present case study in morphological transformation was to highlight the potential of model-based demographic analysis as a means to exploit this wealth of new data. In this Section we thus explicitly demonstrate the advantages of such an approach (as implemented here via SMC ABC) over the standard “one type at a time” mode of study.

Given the (expected) importance of merging as a driver of both star formation and morphological transformation under the canonical hierarchical clustering paradigm (White & Rees 1978) it is perhaps no great surprise that a total of four separate teams have recently attempted con-

straint of the $z \gtrsim 1.5$ (major) merger rate based on close pair and/or asymmetric galaxy counts in deep near-IR imaging. Namely, Bluck et al. (2009, 2011) in the GNS field, Man et al. (2012) in the COSMOS field (in the sub-region of HST/NICMOS coverage), Law et al. (2012) in a dedicated HST/WFC3 study from Cycle 17 sampling multiple fields, and Williams et al. (2011) in the UDS field; the first two exploring the same $M_{\text{gal}} > 10^{11} M_{\odot}$ regime as in our paper and the latter two probing much further down the mass function.

Having identified their target population of impending mergers (in the case of close pair selection) or recent merger remnants (in the case of asymmetric galaxy selection) each of these teams has then proceeded to estimation of the early Universe merger rate in the following manner. First, they compute the merger fraction as the number of mergers detected divided by the total number of galaxies within the target redshift interval and mass range minus the expected proportion of false detections (arising, for instance, from chance alignments along the line-of-sight).²⁰ Second, they adopt a third-party estimate of the merger visibility timescale based on N-body/hydrodynamical simulations of galaxy-galaxy collisions; e.g., Man et al. (2012) adopt $\tau_{\text{close pair}} \sim 0.4 \pm 0.2$ Gyr from Lotz et al. (2008). And third, they either estimate the comoving volume density of their target galaxy population directly from the total count in the survey volume (as in Man et al. 2012), or adopt again a (hopefully more robust, i.e., less prone to cosmic variance error) third-party estimate from the literature (as in Bluck et al. 2009 who take $\Phi_{>10^{11} M_{\odot}}(z)$ from Drory et al. 2005). The merger rate by volume is then computed simply as $f_m \times \Phi/\tau$.

In the lefthand panel of Figure 11 we compare the resulting estimates of $\lambda_m(t)$ from the close pair studies of Bluck et al. (2009) and Man et al. (2012) (denoted BL09 and MA12 here, respectively) against the credible intervals derived from our SMC ABC analysis of the B11 CANDELS/EGS morphological mix (in the “co-evolution” case of our model). We present the aforementioned pair count estimates—which are all well outside our model-based 1σ credible interval—as upper bounds on the major merger rate, since we believe their use of a 1:4 mass ratio threshold may admit a significant number of “weak” accretion events

²⁰ Following Conselice (2006), Bluck et al. (2009, 2011) advocate correction of this raw merger fraction, which they denote f_m , to a “galaxy merger fraction”, denoted f_{gm} , via the relation, $f_{gm} = \frac{2f_m}{1+f_m}$. The motivation for this correction is to identify the “number of galaxies merging as opposed to the number of mergers”. However, we believe this to be a false move in context of their analysis of the most massive galaxies in the GNS, in which they identify impending mergers as those $M_{\text{gal}} > 10^{11} M_{\odot}$ systems host to a close companion (< 30 kpc) of similar brightness (i.e., within ± 1.5 mag in the observed H -band; corresponding to a lower bound on the mass ratio of $\sim 1:4$). Owing to the steepness of the galaxy stellar mass function at $z \gtrsim 1.5$ the vast majority of such companions will almost certainly be sub- $10^{11} M_{\odot}$ —indeed in our sample we find just one close pair in which both are truly high mass galaxies, and Man et al. (2012) find just two. Hence, by “correcting” to f_{gm} as above one is in fact estimating the merger rate experienced by both massive galaxies and the (ill-defined) population of less massive galaxies that will ultimately merge with them, which does not seem a particularly useful exercise.

unlikely to generate a substantial change in morphology; e.g. Hopkins et al. (2009a,b) favour instead a 1:3 mass ratio for the major/minor distinction. Interestingly, the Man et al. (2012) results are at least qualitatively consistent with our posterior inference for the location of the break epoch at $t_{\text{br}} \sim 2.55$ Gyr (since $z = 6$; i.e., $z \sim 1.75$)—though we note in any case that the plotted datapoints do carry rather large uncertainties (of $\sim 88\%$ at the 1σ level; Man et al. 2012), owing in particular to the contribution from cosmic variance error (and necessarily stellar mass estimation error; cf. Brammer et al. 2011 and Section 3.1) in the determination of $\Phi_{>10^{11} M_{\odot}}(z)$.

As a more faithful comparison we present in the righthand panel of Figure 11 the merger rate inferred from application of the $f_m \times \Phi/\tau$ formula to both the asymmetric galaxy counts of Bluck et al. (2011) (denoted here as BL11) and the raw Type IV counts of our B11 CANDELS/EGS dataset (denoted C12). For the former we adopt $\tau_A \approx 0.6 \pm 0.3$ Gyr from Conselice et al. (2009a) (see also Lotz et al. 2008) and for the latter our prior of $\tau_{\text{irr morph}} \approx 0.55 \pm 0.25$ Gyr; and in both cases we employ our in-house estimate of $\Phi_{>10^{11} M_{\odot}}(z)$, which has been calibrated against a large compilation of recent datapoints from the literature in a manner accounting for the presence of systematic biases in the underlying SED-fit based stellar masses of each contributing survey (see Section 3.1 and Figure 2). Given our definition of $\lambda_m(t)$ as the rate of mergers experienced by those systems already in excess of $10^{11} M_{\odot}$ at the stated epoch we must allow for the presence of galaxies only promoted above this mass threshold by the aforesaid accretion event by scaling back our raw count-based estimates according to the factor, $\frac{1}{1+W}$, with W the “nuisance parameter” introduced in Section 3.1. The error bars accompanying each datapoint in this righthand panel of Figure 11 denote the 1σ credible intervals accounting for the relevant uncertainties in these estimates (including, of course, those on the estimation of the population proportion from binomial count data, often mis-handled in astronomical studies; cf. Cameron 2011a).

Immediately evident from inspection of this righthand panel of Figure 11 is the reasonable agreement (well inside the 1σ errors) between our Type IV count-based estimates and those based on the asymmetric galaxy counts from Bluck et al. 2011, confirming a fair degree of equivalence between the use of visual classification and non-parametric, quantitative indicators for the selection of high redshift mergers. Perhaps the most striking impression made by this comparison, however, is the marked offset between the median curve of our SMC ABC constraint on $\lambda_m(t)$ from full demographic analysis of the B11 (CANDELS/EGS) sample and those estimates based only upon the Type IV count at $t \lesssim 2.5$ Gyr (since $z = 6$; i.e., at $z \gtrsim 2.25$). In particular, our median curve here favours a higher merger rate, effectively splitting the difference against the higher asymmetric galaxy count-based datapoint at $t \sim 1.5$ Gyr (since $z = 6$). The explanation for this offset lies, of course, in the fact that we fit our model not just against the merger fraction but rather the full morphological mix, meaning that the fitted merger rate must not only account for the number of ongoing mergers at a given epoch but also the observed population of evolved systems (i.e., ellipticals and bulge-dominated disks) which must have (or in the latter case, *will very likely* have; see




Figure 11. The evolving merger rate over the first ~ 3 Gyr of massive galaxy formation ($z \sim 6$ to $z \sim 1.5$). Here $\lambda_m(t)$ represents the rate by volume of major (*viz.* morphology-changing) mergers experienced by those systems already more massive than $10^{11} M_\odot$ prior to this accretion event. As in Figure 10 above we plot the (pointwise) 1σ and 3σ credible intervals and median curve deriving from our SMC ABC analysis of the B11 (CANDELS/EGS) morphological mix in dark grey, light grey, and blue, respectively. The close pair count-based estimates of BL09 and MA12 (with the former scaled down by a factor of 2 to remove their “correction” from f_m to f_{gm} ; cf. Footnote 20) are overlaid in the lefthand panel; marked here as loose upper bounds since their threshold of a 1:4 mass-ratio for designation as a “major” merger may be overly generous (e.g. Hopkins et al. 2009a,b favour 1:3 alternatively). As a more faithful comparison in the righthand panel we present asymmetric galaxy count-based estimates derived from BL11 (following application of our in-house calibrations for $\Phi_{>10^{11} M_\odot}(z)$ and W ; see Section 3.1 and Figure 2). Also overlaid here are raw Type IV count-based estimates from our (C12) visual classifications, along with an estimate based on the observed Type I and Type II count in our lowest redshift bin; the arrow on the latter noting the fact that the marked position of this datapoint corresponds to the limiting case of all these mergers having occurred with only just enough time to allow fading of the characteristic post-merger irregular features prior to observation.

Section 4.2 below) undergone a merger in their recent past. As an indication of the contribution of Type I and Type II counts²¹ to our fit of $\lambda_m(t)$ we have also marked in Figure 11 a pseudo- λ_m datapoint estimated by treating these evolved galaxy types as ongoing mergers observed at $t_{\text{obs}} - \tau_{\text{Irr morph}}$. The pseudo- λ_m datapoint thus computed confirms that the past rate of merging was very likely to have been higher than that indicated by our raw Type IV counts. Our counts of Type I and II galaxies also contribute a valuable source of data for constraining $\tau_{\text{Irr morph}}$, which through our ABC analysis we verify is indeed probably close to our prior expectation (i.e., our prior mean of $\tau_{\text{Irr morph}} \approx 0.55 \pm 0.25$ falls well within the 1σ credible interval of our posterior, $\tau_{\text{Irr morph}} \approx 0.53 \pm_{0.12}^{0.13}$). The associated increase in confidence regarding the true value of $\tau_{\text{Irr morph}}$ contributes significantly to the reduced width of our SMC ABC credible intervals on $\lambda_m(t)$ (based on the full demographics) relative to those based only on our Type IV counts.

²¹ In fact, we downweight the count of Type II galaxies in this calculation by $(1 - P_{\text{Sph+D remnant}})$ to acknowledge the (minor) role of secular evolution in early bulge formation, as discussed further in Section 4.2 below.

4.2 The Dominant Role of Merging over Secular Evolution for Early Bulge Formation

Another interesting feature of the ABC-based model constraints derived in this paper concerns the relative dominance of merging over secular evolution as the favoured mechanism responsible for building up the first generation of massive bulges in early-type disks. In particular, the posterior probability density of our $P_{\text{Sph+D remnant}}$ parameter favours production of a Type II system in $\approx 33 \pm 17\%$ [1σ] of mergers at these high redshifts—where the gas-rich nature of the progenitors has previously been argued as conducive to disk survival and/or rapid reformation around a central spheroid on the basis of hydrodynamical simulations (e.g. Robertson et al. 2006; Hopkins et al. 2009a; but see Bournaud et al. 2011 regarding the difficulties of reproducing such merger outcomes in models with a realistically cold, turbulent interstellar medium). Interestingly, Hopkins et al. (2009b) have also argued for Type I suppression in gas-rich mergers as a solution to the inconsistency between conventional SAMs and the observed bulge and disk demographics in the local Universe. The posterior density on the secular evolution timescale for our model, $\tau_{\text{sec ev}} \approx 2.8 \pm_{1.2}^{3.8}$ [1σ] Gyr, on the other hand, renders this rival pathway to bulge formation rather unlikely for $z \gtrsim 1.5$ disks. Indeed given the above posterior means one can confirm via repeated simulation from our model that on average only $\sim 7\%$ of Type II systems detected at these redshifts will have been formed via secular evolution. This result is consistent with the recent observations of Genel et al. (2010) and Hopkins et al. (2011) from hydrodynamic simulations in which wind-driven feedback appears to destroy all but the most massive clumps in less time than required for their inwards migration under dynamical friction.

5 CONCLUSIONS

In this paper we have demonstrated the potential of “Approximate Bayesian Computation” for astronomical model analysis through a detailed case study in the morphological transformation of high redshift galaxies. In the process we have derived tight constraints on the evolving merger rate in the early Universe, and revealed the relative dominance of merging over secular evolution for bulge formation at these epochs, through an ABC-based examination of the full population demographics of an $M_{\text{gal}} > 10^{11} M_\odot$, $1.5 < z < 3$ sample from the B11 CANDELS/EGS dataset. More importantly though our exposition of the contemporary “Sequential Monte Carlo” implementation of ABC as well as two modern approaches to summary statistic selection will hopefully guide and inspire further astronomical applications of this powerful statistical technique.

ACKNOWLEDGMENTS

[1] This work is based on observations taken by the CANDELS Multi-Cycle Treasury Program with the NASA/ESA HST, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555.

[2] This study makes use of data from AEGIS, a multiwavelength sky survey conducted with the Chandra, GALEX, Hubble, Keck, CFHT, MMT, Subaru, Palomar, Spitzer, VLA, and other telescopes and supported in part by the NSF, NASA, and the STFC.

[3] E.C. is grateful for financial support from the Australian Research Council.

REFERENCES

- Almeida C., Baugh C.M., Lacey C.G., 2007, *MNRAS*, 376, 1711
- Baldry I.K., Balogh M.L., Bower R.G., Glazebrook K., Nichol R.C., Bamford S.P., Budvari T., 2006, *MNRAS*, 373, 469
- Barnes C., Filippi S., Stumpf M.P.H., Thorne T., 2011, preprint(arXiv:1106.6281)
- Barro G. et al., 2011a, *ApJS*, 193, 13
- Barro G. et al., 2011b, *ApJS*, 193, 30
- Baugh C.M., 2006, *RPPh*, 69, 3101
- Beaumont M.A., Zhang W., Balding D.J., 2002, *Genetics*, 162, 2025
- Bell E.F. et al., 2011, preprint(arXiv:1110.3786)
- Benson A.J., Bower R.G., Frenk C.S., Lacey C.G., Baugh C.M., Cole S., 2003, *ApJ*, 599, 38
- Benson A.J., Borgani S., De Lucia G., Boylan-Kolchin M., Monaco P., 2011, preprint(arXiv:1107.4098)
- Bluck A.F.L., Conselice C.J., Bouwens R.J., Daddi E., Dickinson M., Papovich C., Yan H., 2009, *MNRAS*, 394, 51
- Bluck A.F.L., Conselice C.J., Buitrago F., Grützbauch R., Hoyos C., Mortlock A., Bauer A.E., 2011, preprint(arXiv:1111.5662)
- Bournaud F., Elmegreen B.G., Elmegreen D.M., 2007, *ApJ*, 670, 237
- Bournaud F. et al., 2011, *ApJ*, 730, 4
- Bower R.G., Benson A.J., Malbon R., Helly J.C., Frenk C.S., Baugh C.M., Cole S., Lacey C.G., 2006, *MNRAS*, 370, 645
- Bower R.G., Vernon I., Goldstein M., Benson A.J., Lacey C.G., Baugh C.M., Cole S., Frenk C.S., 2010, *MNRAS*, 407, 2017
- Brammer G.B. et al., 2011, *ApJ*, 739, 24
- Brooks A.M., Governato F., Quinn T., Brook C. B., Wadsley J., 2009, 694, 396
- Buitrago F., Trujillo I., Conselice C.J., Häußler B., 2011, preprint(arXiv:1111.6993)
- Cameron E., Driver S.P., 2009, *A&A*, 493, 489
- Cameron E., 2011, *PASA*, 28, 128
- Cameron E. et al., 2011, *ApJ*, 743, 146
- Capak P.L. et al., 2011, *Nature*, 470, 233
- Caputi K.I., Cirasuolo M., Dunlop J.S., McLure R.J., Farrah D., Almani O., 2010, *MNRAS*, 413, 162
- Carilli C.L., Hodge J., Walter F., Riecher D., Daddi E., Dannerbauer H., Morrison G.E., 2011, *ApJL*, 739, 33
- Chopin N., 2002, *Biometrika*, 89, 539
- Chou R.C.Y., Bridge C.R., Abraham R.G., *AJ*, 141, 87
- Cirasuolo M., McLure R.J., Dunlop J.S., Almaini O., Foucaud S., Simpson C., 2010, *MNRAS*, 401, 1166
- Cole S., Lacey C.G., Baugh C.M., Frenk C.S., 2000, *MNRAS*, 319, 168
- Conselice C.J., Bershadsky M.A., Dickinson M., Papovich C., 2003, *AJ*, 126, 1183
- Conselice C.J., 2006, *ApJ*, 638, 686
- Conselice C.J., Yang C., Bluck A.F.L., 2009, *MNRAS*, 394, 1956
- Conselice C.J., 2009, *MNRAS*, 399, 16
- Conselice C.J. et al., 2011a, *MNRAS*, 417, 2770
- Conselice C.J. et al., 2011b, *MNRAS*, 413, 80
- Csillery K., Blum M.G.B., Gaggiotti F., François O., 2010, *Trends Ecol. Evol.*, 25, 410
- Da Silva R.L., Fumagalli M., Krumholz M., preprint(arXiv:1106.3072)
- Davis M. et al., 2007, *ApJ*, 660, 1
- Dekel A., Sari R., Ceverino D., 2009, *ApJ*, 703, 785
- De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
- De Lucia G., Boylan-Kolchin M., Benson A.J., Fontanot F., Monaco P., 2010, *MNRAS*, 406, 1533
- Del Moral P., Doucet A., Jasra A., 2006, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68, 411
- De Ravel L. et al., 2009, *A&A*, 498, 379
- De Ravel L. et al., 2011, preprint(arXiv:1104.5470)
- Doherty M. et al., 2010, *A&A*, 509, 83
- Driver S.P., Robotham A.S.G., 2010, *MNRAS*, 407, 2131
- Driver S.P. et al., 2006, *MNRAS*, 368, 414
- Drory N., Salvato M., Gabasch A., Bender R., Hopp U., Feulner G., Pannella M., 2005, *ApJL*, 619, 131
- Drovandi C.C., Pettitt A.N., 2010, *Biometrics*, 67, 225
- Drovandi C.C., Pettitt A.N., 2011, *Stat Commun Infect Dis.*, 3, 2
- Elmegreen D.M., Elmegreen B.G., Ravindranath S., Coe D.A., 2007a, *ApJ*, 658, 763
- Elmegreen D.M., Elmegreen B.G., Ferguson T., Mullan B., 2007b, *ApJ*, 663, 734
- Elmegreen B.G., Bournaud F., Elmegreen D.M., 2008, *ApJ*, 688, 67
- Fearnhead P., Prangle D., 2012, *J. R. Statist. Soc. B*, 74, 1
- Fioc M., Rocca-Volmerange B., 1997, *A&A*, 326, 950
- Förster-Schreiber N.M. et al., 2011, *ApJ*, 739, 45
- Fumagalli M., da Silva R.L., Krumholz M.R., 2011, *ApJL*, 741, 26
- Gelman A., Carlin J.B., Stern H.J., Rubin D.B., 2003, *Bayesian Data Analysis (2nd Edition)*, Chapman & Hall/CRC, Boca Raton
- Genzel R. et al., 2011, *ApJ*, 733, 101
- Genel S. et al., 2010, preprint(arXiv:1011.0433)
- Glade N., Ballet P., Bastien O., 2011, preprint(arXiv:1112.1506)
- Goldstein M., Wooff D.A., 2007, *Bayes Linear Statistics: Theory and Methods*, Wiley & Sons, Chichester
- González J.E., Lacey C.G., Baugh C.M., Frenk C.S., Benson A.J., 2009, *MNRAS*, 397, 1254
- Grelaud A., Marin J.-M., Robert C.P., Rodolphe F., Tally F., 2009, *Bayesian Analysis*, 3, 427
- Grogan N. et al., 2011, preprint(arXiv:1105.3753)
- Hand D.J., 1997, *Construction and assessment of classification rules*, Wiley, Chichester (New York)
- Hartig F., Calabrese J.M., Reineking B., Wiegand T., Huth A., 2011, *Ecol. Lett.*, 14, 816
- Hatch N.A., Kurk J.D., Pentericci L., Venemans B.P., Kuiper E., Miley G.K., Röttgering H.J.A., 2011, *MNRAS*, 415, 2993
- Hekker S. et al., 2011, *A&A*, 525, 131

- Henriques B.M.B., Thomas P.A., Oliver S., Roseboom I., 2009, *MNRAS*, 396, 535
- Hopkins P.F., Cox T.J., Younger J.D., Hernquist L., 2009a, *ApJ*, 691, 1168
- Hopkins P.F. et al., 2009b, *MNRAS*, 397, 802
- Hopkins P.F., Kereš D., Murray N., Quataert E., Hernquist L., 2011, preprint(arXiv:1111.6591)
- Ingram A., Done C., 2011, preprint(arXiv:1108.0789)
- Joyce P., Marjoram P., 2008, *Stat. Appl. Genet. Molec. Biol.*, 7, 26
- Kajisawa M. et al., 2009, *ApJ*, 702, 1393
- Kampakoglou M., Trotta R., Silk J., 2008, *MNRAS*, 384, 1414
- Kampczyk P. et al., 2011, preprint(arXiv:1112.4842)
- Kartaltepe J.S. et al., 2011, preprint(arXiv:1110.4057)
- Kelly B.C., Sobolewska M., Seimiginowska A., 2011, *ApJ*, 730, 52
- Kelvin L.S. et al., 2011, preprint(arXiv:1112.1956)
- Kilbinger M. et al., 2010, *MNRAS*, 405, 238
- Kjeldsen H., Bedding T.R., 2004, *Proceedings of the SOHO 14 / GONG 2004 Workshop*, New Haven, p. 101
- Knebe A. et al., 2011, *MNRAS*, 415, 2293
- Kocevski D. et al., 2012, *ApJ*, 744, 148
- Koekemoer A.M. et al., 2011, *ApJS*, 197, 36
- Kramer R.H., Haiman Z., Peng O.S., 2006, *ApJ*, 649, 570
- Lacey C., Cole S., 1993, *MNRAS*, 262, 627
- Law D.R., Steidel C.C., Shapley A.E., Nagy S.R., Reddy N.A., Erb D.K., 2012, *ApJ*, 745, 85
- Lilly S.J., Le Fèvre O., Hammer F., Crampton D., 1996, *ApJL*, 460, 1
- Lin L. et al., 2010, *ApJ*, 718, 1158
- Liu J.S., 2001, *Monte Carlo Strategies in Scientific Computing*, Springer, New York
- López-Sanjuan C. et al., 2009, *ApJ*, 694, 643
- Lotz J.M., Jonsson P., Cox T.J., Primack J.R., *MNRAS*, 391, 1137
- Lotz J.M., Jonsson P., Cox T.J., Primack J.R., 2010, *MNRAS*, 404, 590
- Lotz J.M. et al., 2011, preprint(arXiv:1110.3821)
- Lu Y., Mo H.J., Weinberg M.D., Katz N., 2011, *MNRAS*, 416, 1949
- Madau P., Pozzetti L., Dickinson M.E., Giavalisco M., Steidel C.C., Fruchter A., 1996, *MNRAS*, 283, 138
- Man A., Toft S., Zirm A.W., Wuyts S., van der Wel A., 2012, *ApJ*, 744, 85
- Marjoram P., Tavaré S., 2006, *Genetics*, 7, 759
- Marchesini D. et al., 2010, *ApJ*, 725, 1277
- Marin J.-M., Pillai N., Robert C.P., Rousseau J., 2011, preprint(arXiv:1110.4700)
- Mortlock A. et al., 2011, *MNRAS*, 413, 2845
- Moster B.P., Somerville R.S., Newman J.A., Rix H.-W., 2011, *ApJ*, 731, 113
- Neistein E., Weinmann S.M., 2010, *MNRAS*, 405, 2717
- Newman A.B., Ellis R.S., Bundy K., Treu T., 2011, preprint(arXiv:1110.1637)
- Nott D.J., Fan Y., Sisson S.A., 2011, preprint(arXiv:1112.4755)
- Nunes M.A., Balding D.J., 2010, *Stat. Appl. Genet. Molec. Biol.*, 9, 34
- Oesch P.A. et al., 2010, *ApJL*, 714, 47
- Oesch P.A. et al., 2012, preprint(arXiv:1201.0755)
- Overzier R.A., Guo Q., Kauffmann G., De Lucia G., Bouwens R., Lemson G., 2009, *MNRAS*, 394, 577
- Papovich C. et al., 2011, preprint(arXiv:1110.3794)
- Parkinson H., Cole S., Helly J., 2008, *MNRAS*, 383, 557
- Pipino A., Devriendt J.E.G., Thomas D., Silk J., Kaviraj S., 2009, *A&A*, 505, 1075
- Pritchard J.K., Seielstad M.T., Perez-Lezaun A., Feldman M.W., 1999, *Mol Biol Evol.*, 16, 1791
- Robert C.P., Cornuet J.-M., Marin J.-M., Pillai N.S., 2011, preprint(arXiv:1102.4432)
- Robertson B., Bullock J.S., Cox T.J., Di Matteo T., Hernquist L., Springel V., Yoshida N., 2006, *ApJ*, 645, 986
- Rosario D.J. et al., 2011, preprint(arXiv:1110.3816)
- Schawinski K., Treister E., Urry C.M., Cardamone C.N., Simmons B., Yi S.K., 2011, *ApJL*, 727, 31
- Singh H., Misra N., Hnizdo V., Fedorowicz A., Demchuk E., 2003, *Amer. J. Math. Management Sci.*, 23, 301
- Sisson S., Fan Y., Tanaka M., 2007, *Proc. Natl. Acad. Sci. USA*, 106, 1760
- Skibba R.A. et al., 2009, *MNRAS*, 399, 966
- Spitler L.R. et al., 2011, preprint(arXiv:1112.2691)
- Springel V., White S.D.M., Tormen G., Kauffman G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2005, *Nature*, 435, 629
- Szomoru D. et al., 2010, *ApJL*, 714, 244
- Szomoru D., Franx M., Bouwens R.J., van Dokkum P.G., Labbé I., Illingworth G.D., Trenti M., 2011a, *ApJL*, 735, 22
- Szomoru D., Franx M., van Dokkum P.G., 2011b, preprint(arXiv:1111.3361)
- Tanaka I. et al., 2011, *PASJ*, 63, 415
- Tavaré S., Balding D.J., Griffiths R.C., Donnelly P., 1997, *Genetics*, 145, 505
- Toni T., Stumpf M.P., 2010, *Bioinformatics*, 26, 104
- Tonini C., Maraston C., Ziegler B., Böhm A., Thomas D., Devriendt J., Silk J., 2011, *MNRAS*, 415, 811
- Trenti M., Stiavelli M., 2008, *ApJ*, 676, 767
- Uttley P., McHardy I.M., Papadakis I.E., 2002, *MNRAS*, 332, 231
- van den Bosch F.C., 2000, *ApJ*, 530, 177
- Warren S.J., Hewett P.C., Osmer P.S., 1994, *ApJ*, 421, 412
- Whitaker K.E. et al., 2011, *ApJ*, 735, 86
- White S.D.M., Rees M.J., 1978, *MNRAS*, 183, 341
- Wilkinson R.D., 2008, *Biometrika*, 20, 10
- Williams R.J., Quadri R.F., Franx M., 2011, *ApJL*, 738, 25
- Windhorst R.A. et al., 2011, *ApJS*, 193, 27
- Wood S.N., 2010, *Nature*, 466, 1102
- Wraith D. et al., 2009, *Phys. Rev. D*, 80, 2
- Wuyts S. et al., 2011, *ApJ*, 738, 106

APPENDIX A: LIKELIHOOD COMPUTATION FOR OUR STOCHASTIC MODEL OF MORPHOLOGICAL TRANSFORMATION IN THE “INDEPENDENT EVOLUTION” CASE

In this Appendix we derive the likelihood function, $P(\mathbf{y}|\boldsymbol{\theta})$, of the observed data—i.e., the (HST WFC3/IR) H -band demographics for our sample of 126 galaxies at $1.5 < z < 3$ and $M_{\text{gal}} > 10^{11} M_{\odot}$ selected from the B11 CANDELS/EGS dataset—given a particular set of input parameters, $\boldsymbol{\theta} = \{\alpha_{\text{merge}}, \beta_{\text{merge}}, t_{\text{br}}, P_{\text{Sph+D remnant}}, \tau_{\text{sec ev}}, \tau_{\text{Irr morph}}\}$, under the “independent evolution” case of our stochastic model for high redshift morphological transformation. To this end we must first compute generic expressions for the probability densities of the time of birth and the time of last major merger given an arbitrary redshift of observation. Integration of the conditional transition probabilities of the permitted evolutionary pathways to a given morphology as a function of the above returns the likelihood of observing that type at a particular redshift, and via the independence assumption allows a simple formulation of the complete likelihood function.

Probability Density for the Time of Birth The “birth” of galaxies in our model (i.e., the promotion, via star formation or merging, of new systems to the top end of the $1.5 < z < 6$ stellar mass function) is characterised as a non-homogeneous Poisson process of rate,

$$\lambda_b(t) = 10^K t^\gamma,$$

in units of $\text{Mpc}^{-3}\text{Gyr}^{-1}$ with the origin of the time variable set to $z = 6$. Both K and γ are treated here as nuisance parameters with $f_{K,\gamma} \sim \mathcal{N}_{\text{Trunc}}([-4.1, 0.65]', [0.06^2, 0.1^2]; \rho = 0.05; 0 < \gamma < 1)$. The waiting time distribution for galaxy births under the above-specified Poisson process²² is, of course, exponential in Λ_b -space, where

$$\Lambda_b(t) = \int_0^t \lambda_b(t) dt = \frac{10^K t^{\gamma+1}}{\gamma+1}.$$

Given that each galaxy experiences (by definition) only a single birth the Λ_b -epoch of this event represents a unique draw from the Uniform distribution on $[0, \Lambda_b(t_{\text{obs}})]$ with t_{obs} the cosmological time since $z = 6$ at the observed redshift, z_i . Transformation of variables back to the time domain specifies a probability density for the *time* of birth, t_{birth} , on $[0, t_{\text{obs}}]$ of

$$f_b(t_{\text{birth}}) dt_{\text{birth}} = \frac{\gamma+1}{t_{\text{obs}}^{\gamma+1}} t_{\text{birth}}^\gamma dt_{\text{birth}}.$$

Probability Density for the Time of the Last Major Merger As for the case of the birth function examined above, merging is treated under our stochastic model as a non-homogeneous Poisson process, with a variable rate by volume (in $\text{Mpc}^{-3}\text{Gyr}^{-1}$) set by the input parameters, α_{merge} , β_{merge} , and t_{br} , of

$$\lambda_m(t) = \begin{cases} \frac{\alpha_{\text{merge}} \beta_{\text{merge}}}{t_{\text{br}}^2} t^2 & \text{for } 0 \leq t \leq t_{\text{br}}, \\ \frac{\alpha_{\text{merge}} \beta_{\text{merge}} - (\frac{t-t_{\text{br}}}{t_{1.5}-t_{\text{br}}} \alpha_{\text{merge}} (\beta_{\text{merge}} - 1))}{t_{1.5}-t_{\text{br}}} & \text{for } t_{\text{br}} < t \leq t_{1.5}. \end{cases}$$

²² We note for reference that Glade et al. (2011) provide a brief review of the fundamentals of Poisson processes in their recent paper on the Drake equation.

Here $t_{1.5}$ is used to denote the cosmological time between $z = 6$ and $z = 1.5$ (the lower bound of our sample). The number of mergers experienced by an individual galaxy prior to observation for a given birth time is thus Poisson distributed with rate,

$$\Gamma_m^* = \int_{t_{\text{birth}}}^{t_{\text{obs}}} \frac{\lambda_m(t)}{\Lambda_b(t)} dt = \Gamma_m(t_{\text{obs}}) - \Gamma_m(t_{\text{birth}}),$$

where $\Gamma_m(t) =$

$$\begin{cases} \frac{\alpha_{\text{merge}} \beta_{\text{merge}} (\gamma+1)}{K t_{\text{br}}^2 (2-\gamma)} t^{2-\gamma} & \text{for } 0 \leq t \leq t_{\text{br}}, \\ \begin{cases} \frac{\alpha_{\text{merge}} \beta_{\text{merge}} (\gamma+1)}{K (2-\gamma)} t_{\text{br}}^{-\gamma} + \\ \frac{\alpha_{\text{merge}} \beta_{\text{merge}} (\gamma+1)}{K \gamma} (t_{\text{br}}^{-\gamma} - t^{-\gamma}) - \\ \frac{\alpha_{\text{merge}} (\beta_{\text{merge}} - 1)}{t_{1.5} - t_{\text{br}}} \times \\ \frac{\gamma+1}{K} \left(\frac{t^{1-\gamma}}{1-\gamma} + \frac{t_{\text{br}} t^{-\gamma}}{\gamma} - \frac{t_{\text{br}}^{1-\gamma}}{\gamma(1-\gamma)} \right) \end{cases} & \text{for } t_{\text{br}} < t \leq t_{1.5}. \end{cases}$$

That is, the probability of a galaxy experiencing k mergers between its epoch of birth and its epoch of observation is

$$P(N_m = k | t_{\text{birth}}) = \frac{(\Gamma_m^*)^k e^{-\Gamma_m^*}}{k!},$$

with two cases of particular interest being $P(N_m = 0 | t_{\text{birth}}) = e^{-\Gamma_m^*}$ and $P(N_m > 0 | t_{\text{birth}}) = 1 - e^{-\Gamma_m^*}$. Moreover, for the latter the (non-zero) $N_m = k$ mergers will be distributed uniformly in Γ_m -space on $[\Gamma_m(t_{\text{birth}}), \Gamma_m(t_{\text{obs}})]$. The corresponding probability density of Γ_m -epoch of the most recent merger is then simply that of the k -th order statistic,

$$f_{\Gamma_m, (k)}(\Gamma_m) d\Gamma_m = \frac{k(\Gamma_m - \Gamma_m(t_{\text{birth}}))^{k-1}}{(\Gamma_m^*)^k} d\Gamma_m,$$

(remembering that $\Gamma_m^* = \Gamma_m(t_{\text{obs}}) - \Gamma_m(t_{\text{birth}})$). Summation over all possible (non-zero) merger counts, $k = 1, \dots, \infty$, weighted by their respective probabilities, $P(N_m = k)$, returns the overall probability density of the most recent merger Γ_m -epoch as

$$\begin{aligned} f_{\Gamma_m^r}(\Gamma_m) d\Gamma_m &= \frac{\sum_{k=1}^{\infty} f_{\Gamma_m, (k)}(\Gamma_m) P(N_m = k | t_{\text{birth}})}{P(N_m > 0 | t_{\text{birth}})} d\Gamma_m \\ &= \frac{\sum_{k=1}^{\infty} \frac{k(\Gamma_m - \Gamma_m(t_{\text{birth}}))^{k-1}}{(\Gamma_m^*)^k} \frac{(\Gamma_m^*)^k e^{-\Gamma_m^*}}{k!}}{1 - e^{-\Gamma_m^*}} d\Gamma_m \\ &= \frac{e^{\Gamma_m - \Gamma_m(t_{\text{birth}})}}{e^{\Gamma_m^*} - 1} d\Gamma_m. \end{aligned}$$

Once again transformation of variables delivers the form of this density back in the time domain. Namely,

$$f_{t_m^r}(t_m) dt_m = \frac{e^{\Gamma_m(t_m) - \Gamma_m(t_{\text{birth}})}}{e^{\Gamma_m^*} - 1} \frac{\lambda_m(t_m)}{\Lambda_b(t_m)} dt_m.$$

Likelihood of a Type III Late-Type Disk As described in Section 3.1 galaxies in our model may be born as either Type III disks or Type IV ongoing mergers, with the probability of the latter set by

$$P(C_i(t_{\text{birth}}) = \text{IV} | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) = \frac{W \lambda_m(t_{\text{birth}})}{\lambda_b(t_{\text{birth}})}.$$

We note explicitly here the conditional dependence on our suite of nuisance parameters, $\boldsymbol{\Theta} = \{K, \gamma, W\}$; where this final element, which derives from the shape of the $z \sim 1.5$ – 3 stellar mass function, is taken as $W \approx 0.5 \pm 0.2$, i.e.,

$f_W \sim \mathcal{N}_{\text{Trunc.}}(\mu = 0.5, \sigma = 0.2; W > 0)$. Since no transformation process permits a transition (back) to Type III from outside this state (i.e., Type III represents a transient class of the morphological type Markov chain; see Figure 3) the corresponding likelihood is the easiest to derive—it is merely the probability that a galaxy is born a Type III disk and experiences neither merging nor secular evolution prior to observation.

For a given t_{birth} (with $0 < t_{\text{birth}} < t_{\text{obs}}$) the probability of classification $C_i = \text{III}$ under our stochastic model for morphological transformation is thus

$$\begin{aligned} P_i(C_i = \text{III} | t_{\text{birth}}, \boldsymbol{\theta}) \\ = \iiint_0^\infty P(C_i(t_{\text{birth}}) = \text{III} \cap N_m = 0 \cap S = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) \\ f_{K,\gamma} dK d\gamma f_W dW \\ = \iiint_0^\infty P(C_i(t_{\text{birth}}) = \text{III} | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) P(N_m = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) \\ P(S = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) f_{K,\gamma} dK d\gamma f_W dW. \end{aligned}$$

Here $P(S = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta})$ represents the null secular evolution probability, which operates independently of the null merger probability and the probability of birth as a Type III system. According to the Gamma-distributed form of the secular evolution timescale under our model,

$$P(S = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) = F_{\text{Gamma}(1+50\tau_{\text{sec ev}}, 50)}(t_{\text{obs}} - t_{\text{birth}}),$$

and, as noted above,

$$P(N_m = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) = e^{-\Gamma_m^*}.$$

Integration of this expression by t_{birth} over the corresponding density, $f_b(t_{\text{birth}})$, gives the general likelihood of Type III morphology for the galaxy at hand,

$$P_i(C_i = \text{III} | \boldsymbol{\theta}) = \int_0^{t_{\text{obs}}} P_i(C_i = \text{III} | t_{\text{birth}}, \boldsymbol{\theta}) f_b dt_{\text{birth}}.$$

Analytic solutions for this integral exist in a number of special cases, such as $\gamma = 0$ and $\gamma = 1$ (the latter in terms of the error function), but not to our knowledge for arbitrary, non-integer $\gamma \approx 0.65$ as required to fit the observed build up in number density above $10^{11} M_\odot$ (see Figure 2). To evaluate this expression within our MCMC code we thus employ the efficient numerical technique of Monte Carlo integration (Liu 2001) at run-time.

Likelihood of a Type I Spheroid In contrast to the simple case of the Type III disk presented above there exist an infinite variety of evolutionary pathways potentially leading to the production of a Type I spheroid under our model. However, these pathways may all be considered degenerate with respect to the instance of one final merger (possibly that of the galaxy’s birth) followed by fading of all post-merger irregular features and settling of the merger remnant into a spheroid morphology (see Figure 3). The probability of transition through this penultimate step is, of course, dependent upon the time of the final merger, t_m . Following the derivation above we write down the likelihood of Type I formation for the i -th galaxy with t_{obs} corresponding to z_i as an integral over t_{birth} (and now t_m also) as

$$P_i(C_i = \text{I} | \boldsymbol{\theta}) = \int_0^{t_{\text{obs}}} P_i(C_i = \text{I} | \boldsymbol{\theta}, t_{\text{birth}}) f_b dt_{\text{birth}}$$

with

$$\begin{aligned} P_i(C_i = \text{I} | \boldsymbol{\theta}, t_{\text{birth}}) = \iiint_0^\infty (P(N_m > 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) \\ \int_{t_{\text{birth}}}^{t_{\text{obs}}} P(R = 1 | t_m, \boldsymbol{\theta}, \boldsymbol{\Theta}) P(E = 1 | \boldsymbol{\theta}) f_{t_m^r} dt_m \\ + P(C_i(t_{\text{birth}}) = \text{IV} | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) P(N_m = 0 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) \\ P(R = 1 | t_{\text{birth}}, \boldsymbol{\theta}, \boldsymbol{\Theta}) P(E = 1 | \boldsymbol{\theta})) f_{K,\gamma} dK d\gamma f_W dW. \end{aligned}$$

Here $P(R = 1 | t_m, \boldsymbol{\theta}, \boldsymbol{\Theta}) = F_{\text{Gamma}(1+100\tau_{\text{Irr morph}}, 100)}(t_{\text{obs}} - t_m)$ represents the probability of the post-merger irregular features fading to reveal the final remnant morphology prior to observation, and $P(E = 1 | \boldsymbol{\theta}) = 1 - P_{\text{Sph+D remnant}}$ represents the probability that the final remnant emerges a Type I pure spheroid rather than a Type II spheroid-plus-disk. The first half of this equation corresponds to the case of Type I classification at t_{obs} for a galaxy which has experienced at least one merger since birth, while the second corresponds to the case of a sole, primal merger. Expansion of this expression in full returns another integral with no simple closed form, so once again Monte Carlo integration is ultimately required for its evaluation.

Likelihood of a Type IV Ongoing Merger The likelihood of observing a Type IV ongoing merger is easily derived from the case of the Type I spheroid above with the trivial changes required being a complementation of the probability of the post-merger irregular features fading and removal of the binomial merger remnant type probability.

Likelihood of a Type II Spheroid-plus-Disk As illustrated in Figure 3 there are in fact two non-degenerate pathways to formation of a Type II spheroid-plus-disk system under our model for high redshift morphological transformation, corresponding to merging or secular evolution alternately. So, in principle, writing down the likelihood of this type, $P_i(C_i = \text{II} | \boldsymbol{\theta})$, should be the most difficult of all. However, having derived likelihoods for each of the other morphological types one may simply evaluate this case as

$$\begin{aligned} P_i(C_i = \text{II} | \boldsymbol{\theta}) = 1 - P_i(C_i = \text{I} | \boldsymbol{\theta}) - P_i(C_i = \text{III} | \boldsymbol{\theta}) \\ - P_i(C_i = \text{IV} | \boldsymbol{\theta}). \end{aligned}$$

Such an approach of course requires accurate evaluation of the likelihoods for all three alternative morphologies. Testament to the robustness of our Monte Carlo integration approach is the fact that upon setting $P_{\text{Sph+D remnant}} = 0.5$ and $\tau_{\text{sec ev}} = 100$ Gyr, for which theoretically, $P_i(C_i = \text{I} | \boldsymbol{\theta}) = P_i(C_i = \text{II} | \boldsymbol{\theta})$, the computational agreement of these likelihoods (for a single galaxy) evaluated in R typically extends to at least the 4th significant figure (with just $n_{\text{mc}} \sim 1000$).

Likelihood of the Full Dataset Recovering the likelihood of the full observational dataset, $P(\mathbf{y}, \boldsymbol{\theta})$ with $\mathbf{y} = \{\mathbf{y}_i : (C_i, z_i)\}$ ($i = 1, \dots, N_{\text{gal}}$), in the “independent evolution” case is then simply a matter of taking the product of likelihoods for each individual galaxy (as even neighbouring galaxies evolve uncoupled in this scenario). Thus,

$$P(\mathbf{y} | \boldsymbol{\theta}) = \prod_{i=1}^{N_{\text{gal}}} P_i(C_i | \boldsymbol{\theta}).$$

Of course, when the assumption of independence is relaxed (in order to build a more physically realistic model) in the manner of the “co-evolution” case considered in Section 3.1.1.2 this full dataset likelihood function is no longer so readily tractable, and ABC methods thus become essential for reconstructing the posterior probability densities of the key input parameters.

Markov Chain Monte Carlo Simulation Using the above expressions for $P(\mathbf{y}|\boldsymbol{\theta})$ one may easily generate the benchmark (“tractable likelihood”) posteriors shown in Figure 4 via the (standard) random walk MCMC algorithm. That is, from the current state, $\boldsymbol{\theta}_i$, propose a new state, $\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i + \boldsymbol{\delta}$, by sampling $\boldsymbol{\delta}$ from some zero mean distribution, and accept this proposed state with probability

$$\max\left(\frac{\pi(\boldsymbol{\theta}_{i+1})P(\mathbf{y}|\boldsymbol{\theta}_{i+1})}{\pi(\boldsymbol{\theta}_i)P(\mathbf{y}|\boldsymbol{\theta}_i)}, 1\right).$$

Here we employ the symmetric multivariate Normal distribution, $\mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, for $\boldsymbol{\delta}$ with $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_{\text{prior}}/5$ chosen (by trial-and-error) to produce, on average, an $\sim 40\%$ acceptance rate for $\boldsymbol{\theta}_{i+1}$. Running two separate threads of R on a 2.7 GHz dual core (4 GB ram), 13 inch Macbook Pro laptop we were able to verify satisfactory convergence of this chain after completing a target of 100,000 MCMC steps (with a 1,000 step burn-in period on each thread) in a little less than 48 hours.